

# Welcome to the libuv documentation

## Overview

libuv is a multi-platform support library with a focus on asynchronous I/O. It was primarily developed for use by [Node.js](https://nodejs.org/) [https://nodejs.org/], but it's also used by [Luvit](https://luvit.io/) [https://luvit.io/], [Julia](https://julialang.org/) [https://julialang.org/], [uvloop](https://github.com/MagicStack/uvloop) [https://github.com/MagicStack/uvloop], and [others](https://github.com/libuv/libuv/blob/v1.x/LINKS.md) [https://github.com/libuv/libuv/blob/v1.x/LINKS.md].

## Note

In case you find errors in this documentation you can help by sending [pull requests](https://github.com/libuv/libuv) [https://github.com/libuv/libuv]!

## Features

- Full-featured event loop backed by epoll, kqueue, IOCP, event ports.
- Asynchronous TCP and UDP sockets
- Asynchronous DNS resolution
- Asynchronous file and file system operations
- File system events
- ANSI escape code controlled TTY
- IPC with socket sharing, using Unix domain sockets or named pipes (Windows)
- Child processes
- Thread pool

- Signal handling
- High resolution clock
- Threading and synchronization primitives

## Documentation

- [Design overview](#)
- [API documentation](#)
- [User guide](#)
- [Upgrading](#)

## Downloads

libuv can be downloaded from [here](https://dist.libuv.org/dist/) [https://dist.libuv.org/dist/].

## Installation

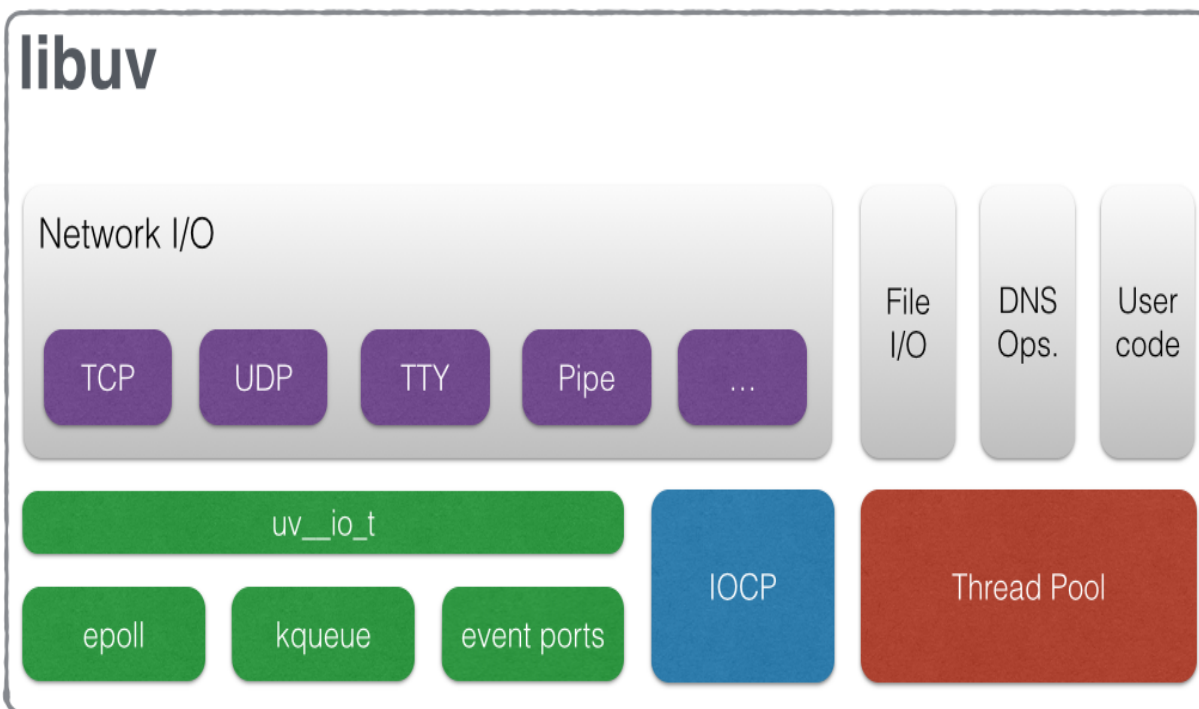
Installation instructions can be found in [the README](https://github.com/libuv/libuv/blob/master/README.md) [https://github.com/libuv/libuv/blob/master/README.md].

# Design overview

libuv is cross-platform support library which was originally written for [Node.js](https://nodejs.org) [https://nodejs.org]. It's designed around the event-driven asynchronous I/O model.

The library provides much more than a simple abstraction over different I/O polling mechanisms: 'handles' and 'streams' provide a high level abstraction for sockets and other entities; cross-platform file I/O and threading functionality is also provided, amongst other things.

Here is a diagram illustrating the different parts that compose libuv and what subsystem they relate to:



## Handles and requests

libuv provides users with 2 abstractions to work with, in combination with the event loop: handles and requests.

Handles represent long-lived objects capable of performing certain operations while active. Some examples:

- A prepare handle gets its callback called once every loop iteration when active.
- A TCP server handle that gets its connection callback called every time there is a new connection.

Requests represent (typically) short-lived operations. These operations can be performed over a handle: write requests are used to write data on a handle; or standalone: getaddrinfo requests don't need a handle they run directly on the loop.

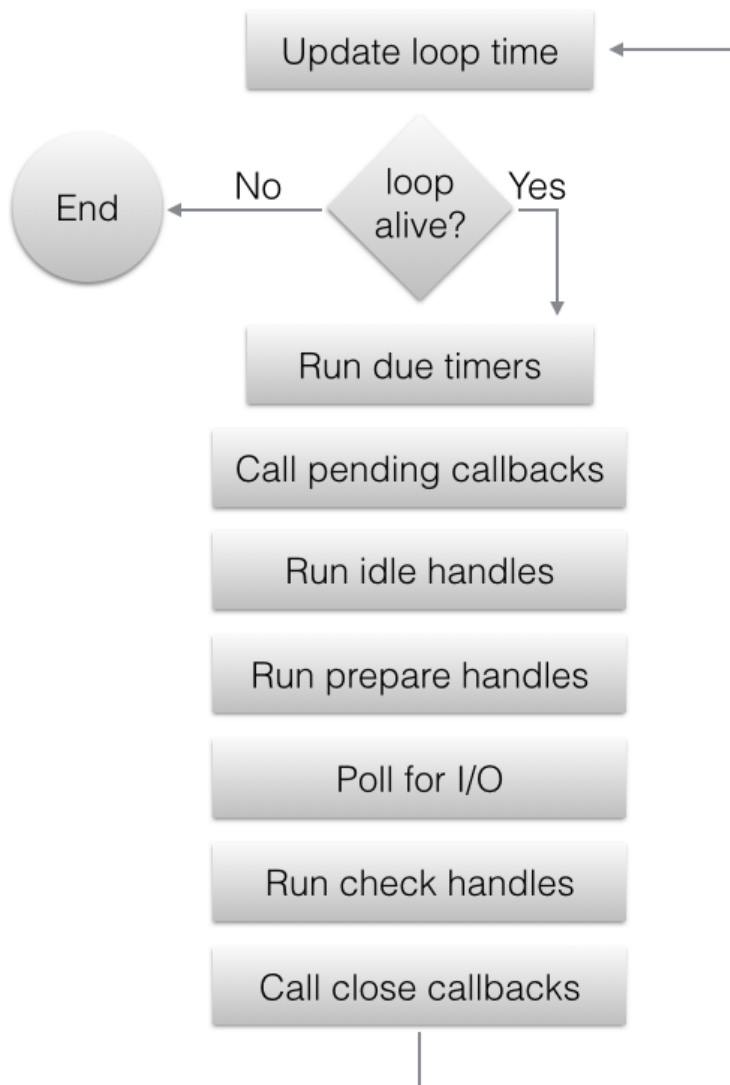
## The I/O loop

The I/O (or event) loop is the central part of libuv. It establishes the content for all I/O operations, and it's meant to be tied to a single thread. One can run multiple event loops as long as each runs in a different thread. The libuv event loop (or any other API involving the loop or handles, for that matter) **is not thread-safe** except where stated otherwise.

The event loop follows the rather usual single threaded asynchronous I/O approach: all (network) I/O is performed on non-blocking sockets which are polled using the best mechanism available on the given platform: epoll on Linux, kqueue on OSX and other BSDs, event ports on SunOS and IOCP on Windows. As part of a loop iteration the loop will block waiting for I/O activity on sockets which have been added to the poller and callbacks will be fired indicating

socket conditions (readable, writable hangup) so handles can read, write or perform the desired I/O operation.

In order to better understand how the event loop operates, the following diagram illustrates all stages of a loop iteration:



1. The loop concept of 'now' is updated. The event loop caches the current time at the start of the event loop tick in order to reduce the number of time-related system calls.

2. If the loop is *alive* an iteration is started, otherwise the loop will exit immediately. So, when is a loop considered to be *alive*? If a loop has active and ref'd handles, active requests or closing handles it's considered to be *alive*.
3. Due timers are run. All active timers scheduled for a time before the loop's concept of *now* get their callbacks called.
4. Pending callbacks are called. All I/O callbacks are called right after polling for I/O, for the most part. There are cases, however, in which calling such a callback is deferred for the next loop iteration. If the previous iteration deferred any I/O callback it will be run at this point.
5. Idle handle callbacks are called. Despite the unfortunate name, idle handles are run on every loop iteration, if they are active.
6. Prepare handle callbacks are called. Prepare handles get their callbacks called right before the loop will block for I/O.
7. Poll timeout is calculated. Before blocking for I/O the loop calculates for how long it should block. These are the rules when calculating the timeout:
  - If the loop was run with the `UV_RUN_NOWAIT` flag, the timeout is 0.
  - If the loop is going to be stopped ([uv\\_stop\(\)](#) was called), the timeout is 0.
  - If there are no active handles or requests, the timeout is 0.
  - If there are any idle handles active, the timeout is 0.

- If there are any handles pending to be closed, the timeout is 0.
  - If none of the above cases matches, the timeout of the closest timer is taken, or if there are no active timers, infinity.
8. The loop blocks for I/O. At this point the loop will block for I/O for the duration calculated in the previous step. All I/O related handles that were monitoring a given file descriptor for a read or write operation get their callbacks called at this point.
  9. Check handle callbacks are called. Check handles get their callbacks called right after the loop has blocked for I/O. Check handles are essentially the counterpart of prepare handles.
  10. Close callbacks are called. If a handle was closed by calling [uv\\_close\(\)](#) it will get the close callback called.
  11. Special case in case the loop was run with UV\_RUN\_ONCE, as it implies forward progress. It's possible that no I/O callbacks were fired after blocking for I/O, but some time has passed so there might be timers which are due, those timers get their callbacks called.
  12. Iteration ends. If the loop was run with UV\_RUN\_NOWAIT or UV\_RUN\_ONCE modes the iteration ends and [uv\\_run\(\)](#) will return. If the loop was run with UV\_RUN\_DEFAULT it will continue from the start if it's still *alive*, otherwise it will also end.

## Important

libuv uses a thread pool to make asynchronous file I/O operations possible, but network I/O is **always** performed

in a single thread, each loop's thread.

## Note

While the polling mechanism is different, libuv makes the execution model consistent across Unix systems and Windows.

## File I/O

Unlike network I/O, there are no platform-specific file I/O primitives libuv could rely on, so the current approach is to run blocking file I/O operations in a thread pool.

For a thorough explanation of the cross-platform file I/O landscape, check out [this post](#)

[<https://blog.libtorrent.org/2012/10/asynchronous-disk-io/>].

libuv currently uses a global thread pool on which all loops can queue work. 3 types of operations are currently run on this pool:

- File system operations
- DNS functions (getaddrinfo and getnameinfo)
- User specified code via [uv\\_queue\\_work\(\)](#).

## Warning

See the [Thread pool work scheduling](#) section for more details, but keep in mind the thread pool size is quite limited.



# API documentation

- [Error handling](#)
- [Version-checking macros and functions](#)
- [uv\\_loop\\_t — Event loop](#)
- [uv\\_handle\\_t — Base handle](#)
- [uv\\_req\\_t — Base request](#)
- [uv\\_timer\\_t — Timer handle](#)
- [uv\\_prepare\\_t — Prepare handle](#)
- [uv\\_check\\_t — Check handle](#)
- [uv\\_idle\\_t — Idle handle](#)
- [uv\\_async\\_t — Async handle](#)
- [uv\\_poll\\_t — Poll handle](#)
- [uv\\_signal\\_t — Signal handle](#)
- [uv\\_process\\_t — Process handle](#)
- [uv\\_stream\\_t — Stream handle](#)
- [uv\\_tcp\\_t — TCP handle](#)
- [uv\\_pipe\\_t — Pipe handle](#)
- [uv\\_tty\\_t — TTY handle](#)
- [uv\\_udp\\_t — UDP handle](#)
- [uv\\_fs\\_event\\_t — FS Event handle](#)
- [uv\\_fs\\_poll\\_t — FS Poll handle](#)
- [File system operations](#)
- [Thread pool work scheduling](#)
- [DNS utility functions](#)
- [Shared library handling](#)
- [Threading and synchronization utilities](#)
- [Miscellaneous utilities](#)
- [Metrics operations](#)

# Error handling

In libuv errors are negative numbered constants. As a rule of thumb, whenever there is a status parameter, or an API functions returns an integer, a negative number will imply an error.

When a function which takes a callback returns an error, the callback will never be called.

## Note

Implementation detail: on Unix error codes are the negated *errno* (or *-errno*), while on Windows they are defined by libuv to arbitrary negative numbers.

## Error constants

### UV\_E2BIG

argument list too long

### UV\_EACCES

permission denied

### UV\_EADDRINUSE

address already in use

### UV\_EADDRNOTAVAIL

address not available

### UV\_EAFNOSUPPORT

address family not supported

**UV\_EAGAIN**

resource temporarily unavailable

**UV\_EAI\_ADDRFAMILY**

address family not supported

**UV\_EAI\_AGAIN**

temporary failure

**UV\_EAI\_BADFLAGS**

bad ai\_flags value

**UV\_EAI\_BADHINTS**

invalid value for hints

**UV\_EAI\_CANCELED**

request canceled

**UV\_EAI\_FAIL**

permanent failure

**UV\_EAI\_FAMILY**

ai\_family not supported

**UV\_EAI\_MEMORY**

out of memory

**UV\_EAI\_NODATA**

no address

**UV\_EAI\_NONAME**

unknown node or service

**UV\_EAI\_OVERFLOW**

argument buffer overflow

**UV\_EAI\_PROTOCOL**

resolved protocol is unknown

**UV\_EAI\_SERVICE**

service not available for socket type

**UV\_EAI\_SOCKTYPE**

socket type not supported

**UV\_EALREADY**

connection already in progress

**UV\_EBADF**

bad file descriptor

**UV\_EBUSY**

resource busy or locked

**UV\_ECANCELED**

operation canceled

**UV\_ECHARSET**

invalid Unicode character

**UV\_ECONNABORTED**

software caused connection abort

**UV\_ECONNREFUSED**

connection refused

**UV\_ECONNRESET**

connection reset by peer

**UV\_EDESTADDRREQ**

destination address required

**UV\_EEXIST**

file already exists

**UV\_EFAULT**

bad address in system call argument

**UV\_EFBIG**

file too large

**UV\_EHOSTUNREACH**

host is unreachable

**UV\_EINTR**

interrupted system call

**UV\_EINVAL**

invalid argument

**UV\_EIO**

i/o error

**UV\_EISCONN**

socket is already connected

**UV\_EISDIR**

illegal operation on a directory

**UV\_ELOOP**

too many symbolic links encountered

**UV\_EMFILE**

too many open files

**UV EMSGSIZE**

message too long

**UV\_ENAMETOOLONG**  
name too long

**UV\_ENETDOWN**  
network is down

**UV\_ENETUNREACH**  
network is unreachable

**UV\_ENFILE**  
file table overflow

**UV\_ENOBUFS**  
no buffer space available

**UV\_ENODEV**  
no such device

**UV\_ENOENT**  
no such file or directory

**UV\_ENOMEM**  
not enough memory

**UV\_ENONET**  
machine is not on the network

**UV\_ENOPROTOOPT**  
protocol not available

**UV\_ENOSPC**  
no space left on device

**UV\_ENOSYS**  
function not implemented

**UV\_ENOTCONN**

socket is not connected

**UV\_ENOTDIR**

not a directory

**UV\_ENOEMPTY**

directory not empty

**UV\_ENOTSOCK**

socket operation on non-socket

**UV\_ENOTSUP**

operation not supported on socket

**UV\_EOVERFLOW**

value too large for defined data type

**UV\_EPERM**

operation not permitted

**UV\_EPIPE**

broken pipe

**UV\_EPROTO**

protocol error

**UV\_EPROTONOSUPPORT**

protocol not supported

**UV\_EPROTOTYPE**

protocol wrong type for socket

**UV\_ERANGE**

result too large

**UV\_EROF**

read-only file system

**UV\_ESHUTDOWN**

cannot send after transport endpoint shutdown

**UV\_ESPIPE**

invalid seek

**UV\_ESRCH**

no such process

**UV\_ETIMEDOUT**

connection timed out

**UV\_ETXTBSY**

text file is busy

**UV\_EXDEV**

cross-device link not permitted

**UV\_UNKNOWN**

unknown error

**UV\_EOF**

end of file

**UV\_ENXIO**

no such device or address

**UV\_EMLINK**

too many links

**UV\_ENOTTY**

inappropriate ioctl for device



**UV\_EFTYPE**

inappropriate file type or format

**UV\_EILSEQ**

illegal byte sequence

**UV\_ESOCKTNOSUPPORT**

socket type not supported

## API

**UV\_ERRNO\_MAP(*iter\_macro*)**

Macro that expands to a series of invocations of *iter\_macro* for each of the error constants above. *iter\_macro* is invoked with two arguments: the name of the error constant without the *UV\_* prefix, and the error message string literal.

**const char\* uv\_strerror(int *err*)**

Returns the error message for the given error code. Leaks a few bytes of memory when you call it with an unknown error code.

**char\* uv\_strerror\_r(int *err*, char\* *buf*, size\_t *buflen*)**

Returns the error message for the given error code. The zero-terminated message is stored in the user-supplied buffer *buf* of at most *buflen* bytes.

*New in version 1.22.0.*

**const char\* uv\_err\_name(int *err*)**

Returns the error name for the given error code. Leaks a few bytes of memory when you call it with an unknown error code.

`char* uv_err_name_r(int err, char* buf, size_t buflen)`

Returns the error name for the given error code. The zero-terminated name is stored in the user-supplied buffer *buf* of at most *buflen* bytes.

*New in version 1.22.0.*

`int uv_translate_sys_error(int sys_errno)`

Returns the libuv error code equivalent to the given platform dependent error code: POSIX error codes on Unix (the ones stored in *errno*), and Win32 error codes on Windows (those returned by *GetLastError()* or *WSAGetLastError()*).

If *sys\_errno* is already a libuv error, it is simply returned.

*Changed in version 1.10.0:* function declared public.

# Version-checking macros and functions

Starting with version 1.0.0 libuv follows the [semantic versioning](https://semver.org) [https://semver.org] scheme. This means that new APIs can be introduced throughout the lifetime of a major release. In this section you'll find all macros and functions that will allow you to write or compile code conditionally, in order to work with multiple libuv versions.

## Macros

### **UV\_VERSION\_MAJOR**

libuv version's major number.

### **UV\_VERSION\_MINOR**

libuv version's minor number.

### **UV\_VERSION\_PATCH**

libuv version's patch number.

### **UV\_VERSION\_IS\_RELEASE**

Set to 1 to indicate a release version of libuv, 0 for a development snapshot.

### **UV\_VERSION\_SUFFIX**

libuv version suffix. Certain development releases such as Release Candidates might have a suffix such as "rc".

### **UV\_VERSION\_HEX**

Returns the libuv version packed into a single integer. 8 bits are used for each component, with the patch number

stored in the 8 least significant bits. E.g. for libuv 1.2.3 this would be 0x010203.

*New in version 1.7.0.*

## Functions

unsigned int **uv\_version**(void)

Returns [UV\\_VERSION\\_HEX](#).

const char\* **uv\_version\_string**(void)

Returns the libuv version number as a string. For non-release versions the version suffix is included.

# uv\_loop\_t — Event loop

The event loop is the central part of libuv's functionality. It takes care of polling for i/o and scheduling callbacks to be run based on different sources of events.

## Data types

`uv_loop_t`

Loop data type.

`void (*uv_walk_cb)(uv_handle_t* handle, void* arg)`

Type definition for callback passed to `uv_walk()`.

## Public members

`void* uv_loop_t.data`

Space for user-defined arbitrary data. libuv does not use and does not touch this field.

## API

`int uv_loop_init(uv_loop_t* loop)`

Initializes the given `uv_loop_t` structure.

`int uv_loop_configure(uv_loop_t* loop,  
uv_loop_option option, ...)`

*New in version 1.0.2.*

Set additional loop options. You should normally call this before the first call to `uv_run()` unless mentioned otherwise.

Returns 0 on success or a UV\_E\* error code on failure. Be prepared to handle UV\_ENOSYS; it means the loop option is not supported by the platform.

Supported options:

- UV\_LOOP\_BLOCK\_SIGNAL: Block a signal when polling for new events. The second argument to [uv\\_loop\\_configure\(\)](#) is the signal number.

This operation is currently only implemented for SIGPROF signals, to suppress unnecessary wakeups when using a sampling profiler. Requesting other signals will fail with UV\_EINVAL.

- UV\_METRICS\_IDLE\_TIME: Accumulate the amount of idle time the event loop spends in the event provider.

This option is necessary to use [uv\\_metrics\\_idle\\_time\(\)](#).

*Changed in version 1.39.0:* added the UV\_METRICS\_IDLE\_TIME option.

`int uv_loop_close(uv\_loop\_t\* loop)`

Releases all internal loop resources. Call this function only when the loop has finished executing and all open handles and requests have been closed, or it will return UV\_EBUSY. After this function returns, the user can free the memory allocated for the loop.

`uv\_loop\_t\* uv_default_loop(void)`

Returns the initialized default loop. It may return NULL in case of allocation failure.

This function is just a convenient way for having a global loop throughout an application, the default loop is in no way different than the ones initialized with [uv\\_loop\\_init\(\)](#).

As such, the default loop can (and should) be closed with [uv\\_loop\\_close\(\)](#) so the resources associated with it are freed.

## Warning

This function is not thread safe.

int [uv\\_run](#)([uv\\_loop\\_t](#)\* *loop*, uv\_run\_mode *mode*)

This function runs the event loop. It will act differently depending on the specified mode:

- `UV_RUN_DEFAULT`: Runs the event loop until there are no more active and referenced handles or requests. Returns non-zero if [uv\\_stop\(\)](#) was called and there are still active handles or requests. Returns zero in all other cases.
- `UV_RUN_ONCE`: Poll for i/o once. Note that this function blocks if there are no pending callbacks. Returns zero when done (no active handles or requests left), or non-zero if more callbacks are expected (meaning you should run the event loop again sometime in the future).
- `UV_RUN_NOWAIT`: Poll for i/o once but don't block if there are no pending callbacks. Returns zero if done (no active handles or requests left), or non-zero if more callbacks are expected (meaning you should run the event loop again sometime in the future).

[uv\\_run\(\)](#) is not reentrant. It must not be called from a callback.

int [uv\\_loop\\_alive](#)(const [uv\\_loop\\_t](#)\* *loop*)

Returns non-zero if there are referenced active handles, active requests or closing handles in the loop.

`void uv_stop(uv\_loop\_t* loop)`

Stop the event loop, causing [uv\\_run\(\)](#) to end as soon as possible. This will happen not sooner than the next loop iteration. If this function was called before blocking for i/o, the loop won't block for i/o on this iteration.

`size_t uv_loop_size(void)`

Returns the size of the *uv\_loop\_t* structure. Useful for FFI binding writers who don't want to know the structure layout.

`int uv_backend_fd(const uv\_loop\_t* loop)`

Get backend file descriptor. Only kqueue, epoll and event ports are supported.

This can be used in conjunction with *uv\_run(loop, UV\_RUN\_NOWAIT)* to poll in one thread and run the event loop's callbacks in another see `test/test-embed.c` for an example.

### **Note**

Embedding a kqueue fd in another kqueue pollset doesn't work on all platforms. It's not an error to add the fd but it never generates events.

`int uv_backend_timeout(const uv\_loop\_t* loop)`

Get the poll timeout. The return value is in milliseconds, or -1 for no timeout.

`uint64_t uv_now(const uv\_loop\_t* loop)`

Return the current timestamp in milliseconds. The timestamp is cached at the start of the event loop tick, see [uv\\_update\\_time\(\)](#) for details and rationale.



The timestamp increases monotonically from some arbitrary point in time. Don't make assumptions about the starting point, you will only get disappointed.

### **Note**

Use [uv\\_hrttime\(\)](#) if you need sub-millisecond granularity.

`void uv_update_time(uv\_loop\_t* loop)`

Update the event loop's concept of "now". Libuv caches the current time at the start of the event loop tick in order to reduce the number of time-related system calls.

You won't normally need to call this function unless you have callbacks that block the event loop for longer periods of time, where "longer" is somewhat subjective but probably on the order of a millisecond or more.

`void uv_walk(uv\_loop\_t* loop, uv\_walk\_cb walk_cb, void* arg)`

Walk the list of handles: *walk\_cb* will be executed with the given *arg*.

`int uv_loop_fork(uv\_loop\_t* loop)`

*New in version 1.12.0.*

Reinitialize any kernel state necessary in the child process after a [fork\(2\)](#) [<https://man7.org/linux/man-pages/man2/fork.2.html>] system call.

Previously started watchers will continue to be started in the child process.

It is necessary to explicitly call this function on every event loop created in the parent process that you plan to

continue to use in the child, including the default loop (even if you don't continue to use it in the parent). This function must be called before calling [uv\\_run\(\)](#) or any other API function using the loop in the child. Failure to do so will result in undefined behaviour, possibly including duplicate events delivered to both parent and child or aborting the child process.

When possible, it is preferred to create a new loop in the child process instead of reusing a loop created in the parent. New loops created in the child process after the fork should not use this function.

This function is not implemented on Windows, where it returns `UV_ENOSYS`.

### **Caution**

This function is experimental. It may contain bugs, and is subject to change or removal. API and ABI stability is not guaranteed.

### **Note**

On Mac OS X, if directory FS event handles were in use in the parent process *for any event loop*, the child process will no longer be able to use the most efficient FSEvent implementation. Instead, uses of directory FS event handles in the child will fall back to the same implementation used for files and on other kqueue-based systems.

### **Caution**

On AIX and SunOS, FS event handles that were already started in the parent process at the time of forking will

*not* deliver events in the child process; they must be closed and restarted. On all other platforms, they will continue to work normally without any further intervention.

### **Caution**

Any previous value returned from [uv\\_backend\\_fd\(\)](#) is now invalid. That function must be called again to determine the correct backend file descriptor.

```
void* uv_loop_get_data(const uv\_loop\_t* loop)
```

Returns *loop->data*.

*New in version 1.19.0.*

```
void* uv_loop_set_data(uv\_loop\_t* loop, void* data)
```

Sets *loop->data* to *data*.

*New in version 1.19.0.*

# uv\_handle\_t — Base handle

`uv_handle_t` is the base type for all libuv handle types.

Structures are aligned so that any libuv handle can be cast to `uv_handle_t`. All API functions defined here work with any handle type.

Libuv handles are not movable. Pointers to handle structures passed to functions must remain valid for the duration of the requested operation. Take care when using stack allocated handles.

## Data types

### `uv_handle_t`

The base libuv handle type.

### `uv_any_handle`

Union of all handle types.

`void (*uv_alloc_cb)(uv\_handle\_t\* handle, size\_t suggested_size, uv\_buf\_t\* buf)`

Type definition for callback passed to [uv\\_read\\_start\(\)](#) and [uv\\_udp\\_recv\\_start\(\)](#). The user must allocate memory and fill the supplied [uv\\_buf\\_t](#) structure. If NULL is assigned as the buffer's base or 0 as its length, a `UV_ENOBUFS` error will be triggered in the [uv\\_udp\\_recv\\_cb](#) or the [uv\\_read\\_cb](#) callback.

Each buffer is used only once and the user is responsible for freeing it in the [uv\\_udp\\_recv\\_cb](#) or the [uv\\_read\\_cb](#) callback.

A suggested size (65536 at the moment in most cases) is provided, but it's just an indication, not related in any way to the pending data to be read. The user is free to allocate the amount of memory they decide.

As an example, applications with custom allocation schemes such as using freelists, allocation pools or slab based allocators may decide to use a different size which matches the memory chunks they already have.

Example:

```
static void my_alloc_cb(uv_handle_t* handle, size_t
suggested_size, uv_buf_t* buf) {
    buf->base = malloc(suggested_size);
    buf->len = suggested_size;
}
```

`void (*uv_close_cb)(uv_handle_t* handle)`

Type definition for callback passed to [uv\\_close\(\)](#).

## Public members

[uv\\_loop\\_t](#)\* `uv_handle_t.loop`

Pointer to the [uv\\_loop\\_t](#) the handle is running on.  
Readonly.

`uv_handle_type` `uv_handle_t.type`

The `uv_handle_type`, indicating the type of the underlying handle. Readonly.

`void*` `uv_handle_t.data`

Space for user-defined arbitrary data. libuv does not use this field.

## API

## **UV\_HANDLE\_TYPE\_MAP(iter\_macro)**

Macro that expands to a series of invocations of *iter\_macro* for each of the handle types. *iter\_macro* is invoked with two arguments: the name of the *uv\_handle\_type* element without the *UV\_* prefix, and the name of the corresponding structure type without the *uv\_* prefix and *\_t* suffix.

## **int uv\_is\_active(const [uv\\_handle\\_t](#)\* handle)**

Returns non-zero if the handle is active, zero if it's inactive. What "active" means depends on the type of handle:

- A *uv\_async\_t* handle is always active and cannot be deactivated, except by closing it with *uv\_close()*.
- A *uv\_pipe\_t*, *uv\_tcp\_t*, *uv\_udp\_t*, etc. handle - basically any handle that deals with i/o - is active when it is doing something that involves i/o, like reading, writing, connecting, accepting new connections, etc.
- A *uv\_check\_t*, *uv\_idle\_t*, *uv\_timer\_t*, etc. handle is active when it has been started with a call to *uv\_check\_start()*, *uv\_idle\_start()*, etc.

Rule of thumb: if a handle of type *uv\_foo\_t* has a *uv\_foo\_start()* function, then it's active from the moment that function is called. Likewise, *uv\_foo\_stop()* deactivates the handle again.

## **int uv\_is\_closing(const [uv\\_handle\\_t](#)\* handle)**

Returns non-zero if the handle is closing or closed, zero otherwise.

### **Note**

This function should only be used between the initialization of the handle and the arrival of the close

callback.

void **uv\_close**([uv\\_handle\\_t](#)\* *handle*, [uv\\_close\\_cb](#) *close\_cb*)

Request handle to be closed. *close\_cb* will be called asynchronously after this call. This **MUST** be called on each handle before memory is released. Moreover, the memory can only be released in *close\_cb* or after it has returned.

Handles that wrap file descriptors are closed immediately but *close\_cb* will still be deferred to the next iteration of the event loop. It gives you a chance to free up any resources associated with the handle.

In-progress requests, like `uv_connect_t` or `uv_write_t`, are cancelled and have their callbacks called asynchronously with `status=UV_ECANCELED`.

void **uv\_ref**([uv\\_handle\\_t](#)\* *handle*)

Reference the given handle. References are idempotent, that is, if a handle is already referenced calling this function again will have no effect.

See [Reference counting](#).

void **uv\_unref**([uv\\_handle\\_t](#)\* *handle*)

Un-reference the given handle. References are idempotent, that is, if a handle is not referenced calling this function again will have no effect.

See [Reference counting](#).

int **uv\_has\_ref**(const [uv\\_handle\\_t](#)\* *handle*)

Returns non-zero if the handle referenced, zero otherwise.

See [Reference counting](#).

`size_t uv_handle_size(uv_handle_type type)`

Returns the size of the given handle type. Useful for FFI binding writers who don't want to know the structure layout.

## Miscellaneous API functions

The following API functions take a [uv\\_handle\\_t](#) argument but they work just for some handle types.

`int uv_send_buffer_size(uv\_handle\_t* handle, int* value)`

Gets or sets the size of the send buffer that the operating system uses for the socket.

If *\*value* == 0, then it will set *\*value* to the current send buffer size. If *\*value* > 0 then it will use *\*value* to set the new send buffer size.

On success, zero is returned. On error, a negative result is returned.

This function works for TCP, pipe and UDP handles on Unix and for TCP and UDP handles on Windows.

### Note

Linux will set double the size and return double the size of the original set value.

`int uv_recv_buffer_size(uv\_handle\_t* handle, int* value)`

Gets or sets the size of the receive buffer that the operating system uses for the socket.



If *\*value* == 0, then it will set *\*value* to the current receive buffer size. If *\*value* > 0 then it will use *\*value* to set the new receive buffer size.

On success, zero is returned. On error, a negative result is returned.

This function works for TCP, pipe and UDP handles on Unix and for TCP and UDP handles on Windows.

### **Note**

Linux will set double the size and return double the size of the original set value.

`int uv_fileno(const uv\_handle\_t* handle, uv\_os\_fd\_t* fd)`  
Gets the platform dependent file descriptor equivalent.

The following handles are supported: TCP, pipes, TTY, UDP and poll. Passing any other handle type will fail with `UV_EINVAL`.

If a handle doesn't have an attached file descriptor yet or the handle itself has been closed, this function will return `UV_EBADF`.

### **Warning**

Be very careful when using this function. libuv assumes it's in control of the file descriptor so any change to it may lead to malfunction.

`uv\_loop\_t* uv_handle_get_loop(const uv\_handle\_t* handle)`  
Returns `handle->loop`.

*New in version 1.19.0.*

`void* uv_handle_get_data(const uv\_handle\_t\* handle)`  
Returns `handle->data`.

*New in version 1.19.0.*

`void* uv_handle_set_data(uv\_handle\_t\* handle, void* data)`  
Sets `handle->data` to `data`.

*New in version 1.19.0.*

`uv_handle_type uv_handle_get_type(const uv\_handle\_t\* handle)`  
Returns `handle->type`.

*New in version 1.19.0.*

`const char* uv_handle_type_name(uv_handle_type type)`  
Returns the name for the equivalent struct for a given handle type, e.g. “*pipe*” (as in [uv\\_pipe\\_t](#)) for `UV_NAMED_PIPE`.

If no such handle type exists, this returns `NULL`.

*New in version 1.19.0.*

## Reference counting

The libuv event loop (if run in the default mode) will run until there are no active *and* referenced handles left. The user can force the loop to exit early by unreferencing handles which are active, for example by calling [uv\\_unref\(\)](#) after calling [uv\\_timer\\_start\(\)](#).

A handle can be referenced or unreferenced, the refcounting scheme doesn't use a counter, so both operations are idempotent.

All handles are referenced when active by default, see [uv\\_is\\_active\(\)](#) for a more detailed explanation on what being *active* involves.

# uv\_req\_t — Base request

*uv\_req\_t* is the base type for all libuv request types.

Structures are aligned so that any libuv request can be cast to *uv\_req\_t*. All API functions defined here work with any request type.

## Data types

**uv\_req\_t**

The base libuv request structure.

**uv\_any\_req**

Union of all request types.

## Public members

**void\* uv\_req\_t.data**

Space for user-defined arbitrary data. libuv does not use this field.

**uv\_req\_type uv\_req\_t.type**

Indicated the type of request. Readonly.

```
typedef enum {
    UV_UNKNOWN_REQ = 0,
    UV_REQ,
    UV_CONNECT,
    UV_WRITE,
    UV_SHUTDOWN,
    UV_UDP_SEND,
    UV_FS,
    UV_WORK,
    UV_GETADDRINFO,
```

```
UV_GETNAMEINFO,  
UV_REQ_TYPE_MAX,  
} uv_req_type;
```

## API

### UV\_REQ\_TYPE\_MAP(*iter\_macro*)

Macro that expands to a series of invocations of *iter\_macro* for each of the request types. *iter\_macro* is invoked with two arguments: the name of the *uv\_req\_type* element without the *UV\_* prefix, and the name of the corresponding structure type without the *uv\_* prefix and *\_t* suffix.

### int [uv\\_cancel](#)([uv\\_req\\_t](#)\* req)

Cancel a pending request. Fails if the request is executing or has finished executing.

Returns 0 on success, or an error code < 0 on failure.

Only cancellation of [uv\\_fs\\_t](#), [uv\\_getaddrinfo\\_t](#), [uv\\_getnameinfo\\_t](#), [uv\\_random\\_t](#) and [uv\\_work\\_t](#) requests is currently supported.

Cancelled requests have their callbacks invoked some time in the future. It's **not** safe to free the memory associated with the request until the callback is called.

Here is how cancellation is reported to the callback:

- A [uv\\_fs\\_t](#) request has its req->result field set to *UV\_ECANCELED*.
- A [uv\\_work\\_t](#), [uv\\_getaddrinfo\\_t](#), [uv\\_getnameinfo\\_t](#) or [uv\\_random\\_t](#) request has its callback invoked with status == *UV\_ECANCELED*.

`size_t uv_req_size(uv_req_type type)`

Returns the size of the given request type. Useful for FFI binding writers who don't want to know the structure layout.

`void* uv_req_get_data(const uv\_req\_t* req)`

Returns `req->data`.

*New in version 1.19.0.*

`void* uv_req_set_data(uv\_req\_t* req, void* data)`

Sets `req->data` to `data`.

*New in version 1.19.0.*

`uv_req_type uv_req_get_type(const uv\_req\_t* req)`

Returns `req->type`.

*New in version 1.19.0.*

`const char* uv_req_type_name(uv_req_type type)`

Returns the name for the equivalent struct for a given request type, e.g. "connect" (as in [uv\\_connect\\_t](#)) for `UV_CONNECT`.

If no such request type exists, this returns `NULL`.

*New in version 1.19.0.*

# [uv\\_timer\\_t](#) — Timer handle

Timer handles are used to schedule callbacks to be called in the future.

## Data types

[uv\\_timer\\_t](#)

Timer handle type.

`void (*uv\_timer\_cb)(uv\_timer\_t* handle)`

Type definition for callback passed to [uv\\_timer\\_start\(\)](#).

## Public members

N/A

## See also

The [uv\\_handle\\_t](#) members also apply.

## API

`int uv\_timer\_init(uv\_loop\_t* loop, uv\_timer\_t* handle)`

Initialize the handle.

`int uv\_timer\_start(uv\_timer\_t* handle, uv\_timer\_cb cb, uint64_t timeout, uint64_t repeat)`

Start the timer. *timeout* and *repeat* are in milliseconds.

If *timeout* is zero, the callback fires on the next event loop iteration. If *repeat* is non-zero, the callback fires first after *timeout* milliseconds and then repeatedly after *repeat* milliseconds.

### **Note**

Does not update the event loop's concept of "now". See [uv\\_update\\_time\(\)](#) for more information.

If the timer is already active, it is simply updated.

```
int uv_timer_stop(uv\_timer\_t* handle)
```

Stop the timer, the callback will not be called anymore.

```
int uv_timer_again(uv\_timer\_t* handle)
```

Stop the timer, and if it is repeating restart it using the repeat value as the timeout. If the timer has never been started before it returns UV\_EINVAL.

```
void uv_timer_set_repeat(uv\_timer\_t* handle,  
uint64_t repeat)
```

Set the repeat interval value in milliseconds. The timer will be scheduled to run on the given interval, regardless of the callback execution duration, and will follow normal timer semantics in the case of a time-slice overrun.

For example, if a 50ms repeating timer first runs for 17ms, it will be scheduled to run again 33ms later. If other tasks consume more than the 33ms following the first timer callback, then the callback will run as soon as possible.

### **Note**



If the repeat value is set from a timer callback it does not immediately take effect. If the timer was non-repeating before, it will have been stopped. If it was repeating, then the old repeat value will have been used to schedule the next timeout.

`uint64_t uv_timer_get_repeat(const uv\_timer\_t* handle)`

Get the timer repeat value.

`uint64_t uv_timer_get_due_in(const uv\_timer\_t* handle)`

Get the timer due value or 0 if it has expired. The time is relative to [uv\\_now\(\)](#).

*New in version 1.40.0.*

## **See also**

The [uv\\_handle\\_t](#) API functions also apply.

# [uv\\_prepare\\_t](#) — Prepare handle

Prepare handles will run the given callback once per loop iteration, right before polling for i/o.

## Data types

`uv_prepare_t`

Prepare handle type.

`void (*uv_prepare_cb)(uv\_prepare\_t\* handle)`

Type definition for callback passed to [uv\\_prepare\\_start\(\)](#).

## Public members

N/A

## See also

The [uv\\_handle\\_t](#) members also apply.

## API

`int uv_prepare_init(uv\_loop\_t\* loop, uv\_prepare\_t\* prepare)`

Initialize the handle. This function always succeeds.

**Returns:** 0

int [uv\\_prepare\\_start](#)([uv\\_prepare\\_t](#)\* prepare,  
[uv\\_prepare\\_cb](#) cb)

Start the handle with the given callback. This function always succeeds, except when *cb* is *NULL*.

**Returns:** 0 on success, or *UV\_EINVAL* when *cb* == *NULL*.

int [uv\\_prepare\\_stop](#)([uv\\_prepare\\_t](#)\* prepare)

Stop the handle, the callback will no longer be called. This function always succeeds.

**Returns:** 0

## See also

The [uv\\_handle\\_t](#) API functions also apply.

# [uv\\_check\\_t](#) — Check handle

Check handles will run the given callback once per loop iteration, right after polling for i/o.

## Data types

[uv\\_check\\_t](#)

Check handle type.

void (\*[uv\\_check\\_cb](#))([uv\\_check\\_t](#)\* *handle*)

Type definition for callback passed to [uv\\_check\\_start\(\)](#).

## Public members

N/A

## See also

The [uv\\_handle\\_t](#) members also apply.

## API

int [uv\\_check\\_init](#)([uv\\_loop\\_t](#)\* *loop*, [uv\\_check\\_t](#)\* *check*)

Initialize the handle. This function always succeeds.

**Returns:** 0

int [uv\\_check\\_start](#)([uv\\_check\\_t](#)\* *check*, [uv\\_check\\_cb](#) *cb*)

Start the handle with the given callback. This function always succeeds, except when *cb* is *NULL*.

**Returns:** 0 on success, or *UV\_EINVAL* when *cb* == *NULL*.

int `uv_check_stop`([uv\\_check\\_t](#)\* *check*)

Stop the handle, the callback will no longer be called.  
This function always succeeds.

**Returns:** 0

### See also

The [uv\\_handle\\_t](#) API functions also apply.

# uv\_idle\_t — Idle handle

Idle handles will run the given callback once per loop iteration, right before the [uv\\_prepare\\_t](#) handles.

## Note

The notable difference with prepare handles is that when there are active idle handles, the loop will perform a zero timeout poll instead of blocking for i/o.

## Warning

Despite the name, idle handles will get their callbacks called on every loop iteration, not when the loop is actually “idle”.

## Data types

`uv_idle_t`

Idle handle type.

`void (*uv_idle_cb)(uv\_idle\_t\* handle)`

Type definition for callback passed to [uv\\_idle\\_start\(\)](#).

## Public members

N/A

## See also

The [uv\\_handle\\_t](#) members also apply.

## API

int [uv\\_idle\\_init](#)([uv\\_loop\\_t](#)\* *loop*, [uv\\_idle\\_t](#)\* *idle*)

Initialize the handle. This function always succeeds.

**Returns:** 0

int [uv\\_idle\\_start](#)([uv\\_idle\\_t](#)\* *idle*, [uv\\_idle\\_cb](#) *cb*)

Start the handle with the given callback. This function always succeeds, except when *cb* is *NULL*.

**Returns:** 0 on success, or *UV\_EINVAL* when *cb* == *NULL*.

int [uv\\_idle\\_stop](#)([uv\\_idle\\_t](#)\* *idle*)

Stop the handle, the callback will no longer be called. This function always succeeds.

**Returns:** 0

### See also

The [uv\\_handle\\_t](#) API functions also apply.

# [uv\\_async\\_t](#) — Async handle

Async handles allow the user to “wakeup” the event loop and get a callback called from another thread.

## Data types

`uv_async_t`

Async handle type.

`void (*uv_async_cb)(uv\_async\_t\* handle)`

Type definition for callback passed to [uv\\_async\\_init\(\)](#).

## Public members

N/A

## See also

The [uv\\_handle\\_t](#) members also apply.

## API

`int uv_async_init(uv\_loop\_t\* loop, uv\_async\_t\* async, uv\_async\_cb async_cb)`

Initialize the handle. A NULL callback is allowed.

**Returns:** 0 on success, or an error code < 0 on failure.

## Note



Unlike other handle initialization functions, it immediately starts the handle.

```
int uv_async_send(uv\_async\_t* async)
```

Wake up the event loop and call the async handle's callback.

**Returns:** 0 on success, or an error code < 0 on failure.

### Note

It's safe to call this function from any thread. The callback will be called on the loop thread.

### Note

[uv\\_async\\_send\(\)](#) is [async-signal-safe](#)

[<https://man7.org/linux/man-pages/man7/signal-safety.7.html>]. It's safe to call this function from a signal handler.

### Warning

libuv will coalesce calls to [uv\\_async\\_send\(\)](#), that is, not every call to it will yield an execution of the callback. For example: if [uv\\_async\\_send\(\)](#) is called 5 times in a row before the callback is called, the callback will only be called once. If [uv\\_async\\_send\(\)](#) is called again after the callback was called, it will be called again.

### See also

The [uv\\_handle\\_t](#) API functions also apply.

## uv\_poll\_t — Poll handle

Poll handles are used to watch file descriptors for readability, writability and disconnection similar to the purpose of [poll\(2\)](https://man7.org/linux/man-pages/man2/poll.2.html). [https://man7.org/linux/man-pages/man2/poll.2.html].

The purpose of poll handles is to enable integrating external libraries that rely on the event loop to signal it about the socket status changes, like c-ares or libssh2. Using `uv_poll_t` for any other purpose is not recommended; [uv\\_tcp\\_t](#), [uv\\_udp\\_t](#), etc. provide an implementation that is faster and more scalable than what can be achieved with [uv\\_poll\\_t](#), especially on Windows.

It is possible that poll handles occasionally signal that a file descriptor is readable or writable even when it isn't. The user should therefore always be prepared to handle EAGAIN or equivalent when it attempts to read from or write to the fd.

It is not okay to have multiple active poll handles for the same socket, this can cause libuv to busyloop or otherwise malfunction.

The user should not close a file descriptor while it is being polled by an active poll handle. This can cause the handle to report an error, but it might also start polling another socket. However the fd can be safely closed immediately after a call to [uv\\_poll\\_stop\(\)](#) or [uv\\_close\(\)](#).

### **Note**

On windows only sockets can be polled with poll handles. On Unix any file descriptor that would be accepted by

[poll\(2\)](https://man7.org/linux/man-pages/man2/poll.2.html) [https://man7.org/linux/man-pages/man2/poll.2.html] can be used.

## Note

On AIX, watching for disconnection is not supported.

## Data types

### `uv_poll_t`

Poll handle type.

`void (*uv_poll_cb)(uv_poll_t* handle, int status, int events)`

Type definition for callback passed to [uv\\_poll\\_start\(\)](#).

### `uv_poll_event`

Poll event types

```
enum uv_poll_event {
    UV_READABLE = 1,
    UV_WRITABLE = 2,
    UV_DISCONNECT = 4,
    UV_PRIORITIZED = 8
};
```

## Public members

N/A

## See also

The [uv\\_handle\\_t](#) members also apply.

# API

`int uv_poll_init(uv\_loop\_t\* loop, uv\_poll\_t\* handle, int fd)`  
Initialize the handle using a file descriptor.

*Changed in version 1.2.2:* the file descriptor is set to non-blocking mode.

`int uv_poll_init_socket(uv\_loop\_t\* loop, uv\_poll\_t\* handle, uv\_os\_sock\_t socket)`

Initialize the handle using a socket descriptor. On Unix this is identical to [uv\\_poll\\_init\(\)](#). On windows it takes a SOCKET handle.

*Changed in version 1.2.2:* the socket is set to non-blocking mode.

`int uv_poll_start(uv\_poll\_t\* handle, int events, uv\_poll\_cb cb)`

Starts polling the file descriptor. *events* is a bitmask made up of `UV_READABLE`, `UV_WRITABLE`, `UV_PRIORITIZED` and `UV_DISCONNECT`. As soon as an event is detected the callback will be called with *status* set to 0, and the detected events set on the *events* field.

The `UV_PRIORITIZED` event is used to watch for sysfs interrupts or TCP out-of-band messages.

The `UV_DISCONNECT` event is optional in the sense that it may not be reported and the user is free to ignore it, but it can help optimize the shutdown path because an extra read or write call might be avoided.

If an error happens while polling, *status* will be  $< 0$  and corresponds with one of the `UV_E*` error codes (see [Error](#)

[handling](#)). The user should not close the socket while the handle is active. If the user does that anyway, the callback *may* be called reporting an error status, but this is **not** guaranteed.

### **Note**

Calling [uv\\_poll\\_start\(\)](#) on a handle that is already active is fine. Doing so will update the events mask that is being watched for.

### **Note**

Though *UV\_DISCONNECT* can be set, it is unsupported on AIX and as such will not be set on the *events* field in the callback.

### **Note**

If one of the events *UV\_READABLE* or *UV\_WRITABLE* are set, the callback will be called again, as long as the given fd/socket remains readable or writable accordingly. Particularly in each of the following scenarios:

- The callback has been called because the socket became readable/writable and the callback did not conduct a read/write on this socket at all.
- The callback committed a read on the socket, and has not read all the available data (when *UV\_READABLE* is set).
- The callback committed a write on the socket, but it remained writable afterwards (when *UV\_WRITABLE* is set).
- The socket has already become readable/writable before calling [uv\\_poll\\_start\(\)](#) on a poll handle

associated with this socket, and since then the state of the socket did not changed.

In all of the above listed scenarios, the socket remains readable or writable and hence the callback will be called again (depending on the events set in the bitmask). This behaviour is known as level triggering.

*Changed in version 1.9.0:* Added the `UV_DISCONNECT` event.

*Changed in version 1.14.0:* Added the `UV_PRIORITIZED` event.

int `uv_poll_stop`([uv\\_poll\\_t](#)\* poll)

Stop polling the file descriptor, the callback will no longer be called.

### **Note**

Calling [uv\\_poll\\_stop\(\)](#) is effective immediately: any pending callback is also canceled, even if the socket state change notification was already pending.

### **See also**

The [uv\\_handle\\_t](#) API functions also apply.

# [uv\\_signal\\_t](#) — Signal handle

Signal handles implement Unix style signal handling on a per-event loop bases.

## Windows notes

Reception of some signals is emulated:

- SIGINT is normally delivered when the user presses CTRL+C. However, like on Unix, it is not generated when terminal raw mode is enabled.
- SIGBREAK is delivered when the user pressed CTRL + BREAK.
- SIGHUP is generated when the user closes the console window. On SIGHUP the program is given approximately 10 seconds to perform cleanup. After that Windows will unconditionally terminate it.
- SIGWINCH is raised whenever libuv detects that the console has been resized. When a libuv app is running under a console emulator, or when a 32-bit libuv app is running on 64-bit system, SIGWINCH will be emulated. In such cases SIGWINCH signals may not always be delivered in a timely manner. For a writable [uv\\_tty\\_t](#) handle libuv will only detect size changes when the cursor is moved. When a readable [uv\\_tty\\_t](#) handle is used, resizing of the console buffer will be detected only if the handle is in raw mode and is being read.
- Watchers for other signals can be successfully created, but these signals are never received. These signals are: *SIGILL*, *SIGABRT*, *SIGFPE*, *SIGSEGV*, *SIGTERM* and *SIGKILL*.

- Calls to `raise()` or `abort()` to programmatically raise a signal are not detected by libuv; these will not trigger a signal watcher.

*Changed in version 1.15.0:* SIGWINCH support on Windows was improved.

*Changed in version 1.31.0:* 32-bit libuv SIGWINCH support on 64-bit Windows was rolled back to old implementation.

## Unix notes

- SIGKILL and SIGSTOP are impossible to catch.
- Handling SIGBUS, SIGFPE, SIGILL or SIGSEGV via libuv results into undefined behavior.
- SIGABRT will not be caught by libuv if generated by `abort()`, e.g. through `assert()`.
- On Linux SIGRT0 and SIGRT1 (signals 32 and 33) are used by the NPTL pthreads library to manage threads. Installing watchers for those signals will lead to unpredictable behavior and is strongly discouraged. Future versions of libuv may simply reject them.

## Data types

`uv_signal_t`

Signal handle type.

`void (*uv_signal_cb)(uv\_signal\_t* handle, int signum)`

Type definition for callback passed to [uv\\_signal\\_start\(\)](#).

## Public members

`int uv_signal_t.signum`

Signal being monitored by this handle. Readonly.



## See also

The [uv\\_handle\\_t](#) members also apply.

## API

int [uv\\_signal\\_init](#)([uv\\_loop\\_t](#)\* loop, [uv\\_signal\\_t](#)\* signal)

Initialize the handle.

int [uv\\_signal\\_start](#)([uv\\_signal\\_t](#)\* signal, [uv\\_signal\\_cb](#) cb, int signum)

Start the handle with the given callback, watching for the given signal.

int [uv\\_signal\\_start\\_oneshot](#)([uv\\_signal\\_t](#)\* signal, [uv\\_signal\\_cb](#) cb, int signum)

*New in version 1.12.0.*

Same functionality as [uv\\_signal\\_start\(\)](#) but the signal handler is reset the moment the signal is received.

int [uv\\_signal\\_stop](#)([uv\\_signal\\_t](#)\* signal)

Stop the handle, the callback will no longer be called.

## See also

The [uv\\_handle\\_t](#) API functions also apply.

# uv\_process\_t — Process handle

Process handles will spawn a new process and allow the user to control it and establish communication channels with it using streams.

## Data types

### uv\_process\_t

Process handle type.

### uv\_process\_options\_t

Options for spawning the process (passed to [uv\\_spawn\(\)](#)).

```
typedef struct uv_process_options_s {
    uv_exit_cb exit_cb;
    const char* file;
    char** args;
    char** env;
    const char* cwd;
    unsigned int flags;
    int stdio_count;
    uv_stdio_container_t* stdio;
    uv_uid_t uid;
    uv_gid_t gid;
} uv_process_options_t;
```

`void (*uv\_exit\_cb)(uv\_process\_t*, int64_t exit_status, int term_signal)`

Type definition for callback passed in [uv\\_process\\_options\\_t](#) which will indicate the exit status and the signal that caused the process to terminate, if any.

## uv\_process\_flags

Flags to be set on the flags field of [uv\\_process\\_options\\_t](#).

```
enum uv_process_flags {
    /*
     * Set the child process' user id.
     */
    UV_PROCESS_SETUID = (1 << 0),
    /*
     * Set the child process' group id.
     */
    UV_PROCESS_SETGID = (1 << 1),
    /*
     * Do not wrap any arguments in quotes, or perform any
other escaping, when
     * converting the argument list into a command line
string. This option is
     * only meaningful on Windows systems. On Unix it is
silently ignored.
     */
    UV_PROCESS_WINDOWS_VERBATIM_ARGUMENTS = (1 << 2),
    /*
     * Spawn the child process in a detached state - this
will make it a process
     * group leader, and will effectively enable the child to
keep running after
     * the parent exits. Note that the child process will
still keep the
     * parent's event loop alive unless the parent process
calls uv_unref() on
     * the child's process handle.
     */
    UV_PROCESS_DETACHED = (1 << 3),
    /*
     * Hide the subprocess window that would normally be
created. This option is
     * only meaningful on Windows systems. On Unix it is
silently ignored.
     */
    UV_PROCESS_WINDOWS_HIDE = (1 << 4),
    /*
     * Hide the subprocess console window that would normally
be created. This
     * option is only meaningful on Windows systems. On Unix
```

```

it is silently
    * ignored.
    */
    UV_PROCESS_WINDOWS_HIDE_CONSOLE = (1 << 5),
    /*
    * Hide the subprocess GUI window that would normally be
    created. This
    * option is only meaningful on Windows systems. On Unix
it is silently
    * ignored.
    */
    UV_PROCESS_WINDOWS_HIDE_GUI = (1 << 6)
};

```

## **uv\_stdio\_container\_t**

Container for each stdio handle or fd passed to a child process.

```

typedef struct uv_stdio_container_s {
    uv_stdio_flags flags;
    union {
        uv_stream_t* stream;
        int fd;
    } data;
} uv_stdio_container_t;

```

## **Public members**

### **int uv\_process\_t.pid**

The PID of the spawned process. It's set after calling [uv\\_spawn\(\)](#).

### **Note**

The [uv\\_handle\\_t](#) members also apply.

### [uv\\_exit\\_cb](#) uv\_process\_options\_t.exit\_cb

Callback called after the process exits.

const char\* **uv\_process\_options\_t.file**

Path pointing to the program to be executed.

char\*\* **uv\_process\_options\_t.args**

Command line arguments. args[0] should be the path to the program. On Windows this uses *CreateProcess* which concatenates the arguments into a string this can cause some strange errors. See the

UV\_PROCESS\_WINDOWS\_VERBATIM\_ARGUMENTS flag on [uv\\_process\\_flags](#).

char\*\* **uv\_process\_options\_t.env**

Environment for the new process. If NULL the parents environment is used.

const char\* **uv\_process\_options\_t.cwd**

Current working directory for the subprocess.

unsigned int **uv\_process\_options\_t.flags**

Various flags that control how [uv\\_spawn\(\)](#) behaves. See [uv\\_process\\_flags](#).

int **uv\_process\_options\_t.stdio\_count**

[uv\\_stdio\\_container\\_t](#)\* **uv\_process\_options\_t.stdio**

The *stdio* field points to an array of [uv\\_stdio\\_container\\_t](#) structs that describe the file descriptors that will be made available to the child process. The convention is that stdio[0] points to stdin, fd 1 is used for stdout, and fd 2 is stderr.

### **Note**

On Windows file descriptors greater than 2 are available to the child process only if the child processes uses the MSVCRT runtime.

`uv_uid_t uv_process_options_t.uid`

`uv_gid_t uv_process_options_t.gid`

Libuv can change the child process' user/group id. This happens only when the appropriate bits are set in the flags fields.

### **Note**

This is not supported on Windows, [uv\\_spawn\(.\)](#) will fail and set the error to UV\_ENOTSUP.

`uv_stdio_flags uv_stdio_container_t.flags`

Flags specifying how the stdio container should be passed to the child.

`union @0 uv_stdio_container_t.data`

Union containing either the *stream* or *fd* to be passed on to the child process.

## **API**

`void uv_disable_stdio_inheritance(void)`

Disables inheritance for file descriptors / handles that this process inherited from its parent. The effect is that child processes spawned by this process don't accidentally inherit these handles.

It is recommended to call this function as early in your program as possible, before the inherited file descriptors can be closed or duplicated.

### **Note**

This function works on a best-effort basis: there is no guarantee that libuv can discover all file descriptors that were inherited. In general it does a better job on Windows than it does on Unix.

`int uv_spawn(uv\_loop\_t\* loop, uv\_process\_t\* handle, const uv\_process\_options\_t\* options)`

Initializes the process handle and starts the process. If the process is successfully spawned, this function will return 0. Otherwise, the negative error code corresponding to the reason it couldn't spawn is returned.

Possible reasons for failing to spawn would include (but not be limited to) the file to execute not existing, not having permissions to use the setuid or setgid specified, or not having enough memory to allocate for the new process.

*Changed in version 1.24.0: Added `UV_PROCESS_WINDOWS_HIDE_CONSOLE` and `UV_PROCESS_WINDOWS_HIDE_GUI` flags.*

`int uv_process_kill(uv\_process\_t\* handle, int signum)`

Sends the specified signal to the given process handle. Check the documentation on [uv\\_signal\\_t — Signal handle](#) for signal support, specially on Windows.

`int uv_kill(int pid, int signum)`

Sends the specified signal to the given PID. Check the documentation on [uv\\_signal\\_t — Signal handle](#) for signal support, specially on Windows.

`uv\_pid\_t uv_process_get_pid(const uv\_process\_t\* handle)`

Returns *handle->pid*.

*New in version 1.19.0.*

### **See also**

The [uv\\_handle\\_t](#) API functions also apply.



# uv\_stream\_t — Stream handle

Stream handles provide an abstraction of a duplex communication channel. [uv\\_stream\\_t](#) is an abstract type, libuv provides 3 stream implementations in the form of [uv\\_tcp\\_t](#), [uv\\_pipe\\_t](#) and [uv\\_tty\\_t](#).

## Data types

**uv\_stream\_t**

Stream handle type.

**uv\_connect\_t**

Connect request type.

**uv\_shutdown\_t**

Shutdown request type.

**uv\_write\_t**

Write request type. Careful attention must be paid when reusing objects of this type. When a stream is in non-blocking mode, write requests sent with `uv_write` will be queued. Reusing objects at this point is undefined behaviour. It is safe to reuse the `uv_write_t` object only after the callback passed to `uv_write` is fired.

`void (*uv_read_cb)(uv_stream_t* stream, ssize_t nread, const uv_buf_t* buf)`

Callback called when data was read on a stream.

*nread* is  $> 0$  if there is data available or  $< 0$  on error. When we've reached EOF, *nread* will be set to `UV_EOF`. When *nread*  $< 0$ , the *buf* parameter might not point to a

valid buffer; in that case *buf.len* and *buf.base* are both set to 0.

### **Note**

*nread* might be 0, which does *not* indicate an error or EOF. This is equivalent to EAGAIN or EWOULDBLOCK under `read(2)`.

The callee is responsible for stopping/closing the stream when an error happens by calling [uv\\_read\\_stop\(\)](#) or [uv\\_close\(\)](#). Trying to read from the stream again is undefined.

The callee is responsible for freeing the buffer, libuv does not reuse it. The buffer may be a null buffer (where *buf->base* == NULL and *buf->len* == 0) on error.

`void (*uv_write_cb)(uv_write_t* req, int status)`

Callback called after data was written on a stream. *status* will be 0 in case of success, < 0 otherwise.

`void (*uv_connect_cb)(uv_connect_t* req, int status)`

Callback called after a connection started by `uv_connect()` is done. *status* will be 0 in case of success, < 0 otherwise.

`void (*uv_shutdown_cb)(uv_shutdown_t* req, int status)`

Callback called after a shutdown request has been completed. *status* will be 0 in case of success, < 0 otherwise.

`void (*uv_connection_cb)(uv_stream_t* server, int status)`

Callback called when a stream server has received an incoming connection. The user can accept the connection

by calling [uv\\_accept\(\)](#). *status* will be 0 in case of success, < 0 otherwise.

## Public members

[size\\_t uv\\_stream\\_t.write\\_queue\\_size](#)

Contains the amount of queued bytes waiting to be sent. Readonly.

[uv\\_stream\\_t\\* uv\\_connect\\_t.handle](#)

Pointer to the stream where this connection request is running.

[uv\\_stream\\_t\\* uv\\_shutdown\\_t.handle](#)

Pointer to the stream where this shutdown request is running.

[uv\\_stream\\_t\\* uv\\_write\\_t.handle](#)

Pointer to the stream where this write request is running.

[uv\\_stream\\_t\\* uv\\_write\\_t.send\\_handle](#)

Pointer to the stream being sent using this write request.

## See also

The [uv\\_handle\\_t](#) members also apply.

## API

`int uv_shutdown(uv\_shutdown\_t\* req, uv\_stream\_t\* handle, uv\_shutdown\_cb cb)`

Shutdown the outgoing (write) side of a duplex stream. It waits for pending write requests to complete. The *handle* should refer to a initialized stream. *req* should be an

uninitialized shutdown request struct. The *cb* is called after shutdown is complete.

```
int uv_listen(uv\_stream\_t* stream, int backlog,  
uv\_connection\_cb cb)
```

Start listening for incoming connections. *backlog* indicates the number of connections the kernel might queue, same as [listen\(2\)](#) [<https://man7.org/linux/man-pages/man2/listen.2.html>]. When a new incoming connection is received the [uv\\_connection\\_cb](#) callback is called.

```
int uv_accept(uv\_stream\_t* server, uv\_stream\_t* client)
```

This call is used in conjunction with [uv\\_listen\(\)](#) to accept incoming connections. Call this function after receiving a [uv\\_connection\\_cb](#) to accept the connection. Before calling this function the client handle must be initialized. < 0 return value indicates an error.

When the [uv\\_connection\\_cb](#) callback is called it is guaranteed that this function will complete successfully the first time. If you attempt to use it more than once, it may fail. It is suggested to only call this function once per [uv\\_connection\\_cb](#) call.

### **Note**

*server* and *client* must be handles running on the same loop.

```
int uv_read_start(uv\_stream\_t* stream,  
uv\_alloc\_cb alloc_cb, uv\_read\_cb read_cb)
```

Read data from an incoming stream. The [uv\\_read\\_cb](#) callback will be made several times until there is no more data to read or [uv\\_read\\_stop\(\)](#) is called.

Changed in version 1.38.0: [uv\\_read\\_start\(\)](#) now consistently returns `UV_EALREADY` when called twice, and `UV_EINVAL` when the stream is closing. With older libuv versions, it returns `UV_EALREADY` on Windows but not UNIX, and `UV_EINVAL` on UNIX but not Windows.

`int uv_read_stop(uv\_stream\_t\*)`

Stop reading data from the stream. The [uv\\_read\\_cb](#) callback will no longer be called.

This function is idempotent and may be safely called on a stopped stream.

This function will always succeed; hence, checking its return value is unnecessary. A non-zero return indicates that finishing releasing resources may be pending on the next input event on that TTY on Windows, and does not indicate failure.

`int uv_write(uv\_write\_t\* req, uv\_stream\_t\* handle, const uv\_buf\_t bufs[], unsigned int nbufs, uv\_write\_cb cb)`

Write data to stream. Buffers are written in order.

Example:

```
void cb(uv_write_t* req, int status) {
    /* Logic which handles the write result */
}

uv_buf_t a[] = {
    { .base = "1", .len = 1 },
    { .base = "2", .len = 1 }
};

uv_buf_t b[] = {
    { .base = "3", .len = 1 },
    { .base = "4", .len = 1 }
};

uv_write_t req1;
```

```
uv_write_t req2;

/* writes "1234" */
uv_write(&req1, stream, a, 2, cb);
uv_write(&req2, stream, b, 2, cb);
```

### Note

The memory pointed to by the buffers must remain valid until the callback gets called. This also holds for [uv\\_write2\(\)](#).

```
int uv_write2(uv\_write\_t* req, uv\_stream\_t* handle, const uv\_buf\_t bufs[], unsigned int nbufs, uv\_stream\_t* send_handle, uv\_write\_cb cb)
```

Extended write function for sending handles over a pipe. The pipe must be initialized with `ipc == 1`.

### Note

`send_handle` must be a TCP, pipe and UDP handle on Unix, or a TCP handle on Windows, which is a server or a connection (listening or connected state). Bound sockets or pipes will be assumed to be servers.

```
int uv_try_write(uv\_stream\_t* handle, const uv\_buf\_t bufs[], unsigned int nbufs)
```

Same as [uv\\_write\(\)](#), but won't queue a write request if it can't be completed immediately.

Will return either:

- > 0: number of bytes written (can be less than the supplied buffer size).
- < 0: negative error code (`UV_EAGAIN` is returned if no data can be sent immediately).

```
int uv_try_write2(uv\_stream\_t* handle, const
uv\_buf\_t bufs[], unsigned int nbufs,
uv\_stream\_t* send_handle)
```

Same as [uv\\_try\\_write\(\)](#) and extended write function for sending handles over a pipe like c:func:uv\_write2.

Try to send a handle is not supported on Windows, where it returns UV\_EAGAIN.

*New in version 1.42.0.*

```
int uv_is_readable(const uv\_stream\_t* handle)
```

Returns 1 if the stream is readable, 0 otherwise.

```
int uv_is_writable(const uv\_stream\_t* handle)
```

Returns 1 if the stream is writable, 0 otherwise.

```
int uv_stream_set_blocking(uv\_stream\_t* handle,
int blocking)
```

Enable or disable blocking mode for a stream.

When blocking mode is enabled all writes complete synchronously. The interface remains unchanged otherwise, e.g. completion or failure of the operation will still be reported through a callback which is made asynchronously.

### **Warning**

Relying too much on this API is not recommended. It is likely to change significantly in the future.

Currently only works on Windows for [uv\\_pipe\\_t](#) handles. On UNIX platforms, all [uv\\_stream\\_t](#) handles are supported.

Also libuv currently makes no ordering guarantee when the blocking mode is changed after write requests have already been submitted. Therefore it is recommended to set the blocking mode immediately after opening or creating the stream.

*Changed in version 1.4.0:* UNIX implementation added.

`size_t uv_stream_get_write_queue_size(const uv\_stream\_t* stream)`

Returns `stream->write_queue_size`.

*New in version 1.19.0.*

## **See also**

The [uv\\_handle\\_t](#) API functions also apply.



# [uv\\_tcp\\_t](#) — TCP handle

TCP handles are used to represent both TCP streams and servers.

[uv\\_tcp\\_t](#) is a 'subclass' of [uv\\_stream\\_t](#).

## Data types

`uv_tcp_t`

TCP handle type.

## Public members

N/A

## See also

The [uv\\_stream\\_t](#) members also apply.

## API

`int uv_tcp_init(uv\_loop\_t* loop, uv\_tcp\_t* handle)`

Initialize the handle. No socket is created as of yet.

`int uv_tcp_init_ex(uv\_loop\_t* loop, uv\_tcp\_t* handle, unsigned int flags)`

Initialize the handle with the specified flags. At the moment only the lower 8 bits of the *flags* parameter are used as the socket domain. A socket will be created for

the given domain. If the specified domain is AF\_UNSPEC no socket is created, just like [uv\\_tcp\\_init\(\)](#).

*New in version 1.7.0.*

int [uv\\_tcp\\_open](#)([uv\\_tcp\\_t](#)\* handle, [uv\\_os\\_sock\\_t](#) sock)

Open an existing file descriptor or SOCKET as a TCP handle.

*Changed in version 1.2.1:* the file descriptor is set to non-blocking mode.

### **Note**

The passed file descriptor or SOCKET is not checked for its type, but it's required that it represents a valid stream socket.

int [uv\\_tcp\\_nodelay](#)([uv\\_tcp\\_t](#)\* handle, int enable)

Enable `TCP_NODELAY`, which disables Nagle's algorithm.

int [uv\\_tcp\\_keepalive](#)([uv\\_tcp\\_t](#)\* handle, int enable, unsigned int delay)

Enable / disable TCP keep-alive. *delay* is the initial delay in seconds, ignored when *enable* is zero.

After *delay* has been reached, 10 successive probes, each spaced 1 second from the previous one, will still happen. If the connection is still lost at the end of this procedure, then the handle is destroyed with a `UV_ETIMEDOUT` error passed to the corresponding callback.

int [uv\\_tcp\\_simultaneous\\_accepts](#)([uv\\_tcp\\_t](#)\* handle, int enable)

Enable / disable simultaneous asynchronous accept requests that are queued by the operating system when listening for new TCP connections.

This setting is used to tune a TCP server for the desired performance. Having simultaneous accepts can significantly improve the rate of accepting connections (which is why it is enabled by default) but may lead to uneven load distribution in multi-process setups.

```
int uv_tcp_bind(uv\_tcp\_t* handle, const struct  
sockaddr* addr, unsigned int flags)
```

Bind the handle to an address and port. *addr* should point to an initialized struct `sockaddr_in` or struct `sockaddr_in6`.

When the port is already taken, you can expect to see an `UV_EADDRINUSE` error from [uv\\_listen\(\)](#) or [uv\\_tcp\\_connect\(\)](#). That is, a successful call to this function does not guarantee that the call to [uv\\_listen\(\)](#) or [uv\\_tcp\\_connect\(\)](#) will succeed as well.

*flags* can contain `UV_TCP_IPV6ONLY`, in which case dual-stack support is disabled and only IPv6 is used.

```
int uv_tcp_getsockname(const uv\_tcp\_t* handle, struct  
sockaddr* name, int* namelen)
```

Get the current address to which the handle is bound. *name* must point to a valid and big enough chunk of memory, struct `sockaddr_storage` is recommended for IPv4 and IPv6 support.

```
int uv_tcp_getpeername(const uv\_tcp\_t* handle, struct  
sockaddr* name, int* namelen)
```

Get the address of the peer connected to the handle. *name* must point to a valid and big enough chunk of memory, `struct sockaddr_storage` is recommended for IPv4 and IPv6 support.

```
int uv_tcp_connect(uv\_connect\_t\* req, uv\_tcp\_t\* handle,  
const struct sockaddr* addr, uv\_connect\_cb cb)
```

Establish an IPv4 or IPv6 TCP connection. Provide an initialized TCP handle and an uninitialized [uv\\_connect\\_t](#). *addr* should point to an initialized `struct sockaddr_in` or `struct sockaddr_in6`.

On Windows if the *addr* is initialized to point to an unspecified address (`0.0.0.0` or `::`) it will be changed to point to `localhost`. This is done to match the behavior of Linux systems.

The callback is made when the connection has been established or when a connection error happened.

*Changed in version 1.19.0:* added `0.0.0.0` and `::` to `localhost` mapping

## See also

The [uv\\_stream\\_t](#) API functions also apply.

```
int uv_tcp_close_reset(uv\_tcp\_t\* handle,  
uv\_close\_cb close_cb)
```

Resets a TCP connection by sending a RST packet. This is accomplished by setting the `SO_LINGER` socket option with a linger interval of zero and then calling [uv\\_close\(\)](#). Due to some platform inconsistencies, mixing of [uv\\_shutdown\(\)](#) and [uv\\_tcp\\_close\\_reset\(\)](#) calls is not allowed.

*New in version 1.32.0.*

int **uv\_socketpair**(int *type*, int *protocol*,  
[uv\\_os\\_sock\\_t](#) *socket\_vector*[2], int *flags0*, int *flags1*)

Create a pair of connected sockets with the specified properties. The resulting handles can be passed to *uv\_tcp\_open*, used with *uv\_spawn*, or for any other purpose.

Valid values for *flags0* and *flags1* are:

- **UV\_NONBLOCK\_PIPE**: Opens the specified socket handle for *OVERLAPPED* or *FIONBIO/O\_NONBLOCK* I/O usage. This is recommended for handles that will be used by libuv, and not usually recommended otherwise.

Equivalent to [socketpair\(2\)](#) [https://man7.org/linux/man-pages/man2/socketpair.2.html] with a domain of **AF\_UNIX**.

*New in version 1.41.0.*

# [uv\\_pipe\\_t](#) — Pipe handle

Pipe handles provide an abstraction over streaming files on Unix (including local domain sockets, pipes, and FIFOs) and named pipes on Windows.

[uv\\_pipe\\_t](#) is a 'subclass' of [uv\\_stream\\_t](#).

## Data types

`uv_pipe_t`

Pipe handle type.

## Public members

`int uv_pipe_t.ipc`

Whether this pipe is suitable for handle passing between processes. Only a connected pipe that will be passing the handles should have this flag set, not the listening pipe that `uv_accept` is called on.

## See also

The [uv\\_stream\\_t](#) members also apply.

## API

`int uv_pipe_init(uv\_loop\_t* loop, uv\_pipe\_t* handle, int ipc)`

Initialize a pipe handle. The *ipc* argument is a boolean to indicate if this pipe will be used for handle passing between processes (which may change the bytes on the

wire). Only a connected pipe that will be passing the handles should have this flag set, not the listening pipe that `uv_accept` is called on.

```
int uv_pipe_open(uv\_pipe\_t\* handle, uv\_file file)
```

Open an existing file descriptor or HANDLE as a pipe.

*Changed in version 1.2.1:* the file descriptor is set to non-blocking mode.

### **Note**

The passed file descriptor or HANDLE is not checked for its type, but it's required that it represents a valid pipe.

```
int uv_pipe_bind(uv\_pipe\_t\* handle, const char* name)
```

Bind the pipe to a file path (Unix) or a name (Windows).

### **Note**

Paths on Unix get truncated to `sizeof(sockaddr_un.sun_path)` bytes, typically between 92 and 108 bytes.

```
void uv_pipe_connect(uv\_connect\_t\* req, uv\_pipe\_t\* handle, const char* name, uv\_connect\_cb cb)
```

Connect to the Unix domain socket or the named pipe.

### **Note**

Paths on Unix get truncated to `sizeof(sockaddr_un.sun_path)` bytes, typically between 92 and 108 bytes.

```
int uv_pipe_getsockname(const uv_pipe_t* handle,  
char* buffer, size_t* size)
```

Get the name of the Unix domain socket or the named pipe.

A preallocated buffer must be provided. The size parameter holds the length of the buffer and it's set to the number of bytes written to the buffer on output. If the buffer is not big enough UV\_ENOBUFS will be returned and len will contain the required size.

*Changed in version 1.3.0:* the returned length no longer includes the terminating null byte, and the buffer is not null terminated.

```
int uv_pipe_getpeername(const uv_pipe_t* handle,  
char* buffer, size_t* size)
```

Get the name of the Unix domain socket or the named pipe to which the handle is connected.

A preallocated buffer must be provided. The size parameter holds the length of the buffer and it's set to the number of bytes written to the buffer on output. If the buffer is not big enough UV\_ENOBUFS will be returned and len will contain the required size.

*New in version 1.3.0.*

```
void uv_pipe_pending_instances(uv_pipe_t* handle,  
int count)
```

Set the number of pending pipe instance handles when the pipe server is waiting for connections.

### **Note**

This setting applies to Windows only.



`int uv_pipe_pending_count(uv\_pipe\_t\* handle)`

`uv_handle_type uv_pipe_pending_type(uv\_pipe\_t\* handle)`

Used to receive handles over IPC pipes.

First - call [uv\\_pipe\\_pending\\_count\(\)](#), if it's > 0 then initialize a handle of the given type, returned by

[uv\\_pipe\\_pending\\_type\(\)](#) and call `uv_accept(pipe, handle)`.

## See also

The [uv\\_stream\\_t](#) API functions also apply.

`int uv_pipe_chmod(uv\_pipe\_t\* handle, int flags)`

Alters pipe permissions, allowing it to be accessed from processes run by different users. Makes the pipe writable or readable by all users. Mode can be `UV_WRITABLE`, `UV_READABLE` or `UV_WRITABLE | UV_READABLE`. This function is blocking.

*New in version 1.16.0.*

`int uv_pipe(uv\_file fds[2], int read_flags, int write_flags)`

Create a pair of connected pipe handles. Data may be written to `fds[1]` and read from `fds[0]`. The resulting handles can be passed to `uv_pipe_open`, used with `uv_spawn`, or for any other purpose.

Valid values for *flags* are:

- `UV_NONBLOCK_PIPE`: Opens the specified socket handle for `OVERLAPPED` or `FIONBIO/O_NONBLOCK` I/O usage. This is

recommended for handles that will be used by libuv, and not usually recommended otherwise.

Equivalent to [pipe\(2\)](https://man7.org/linux/man-pages/man2/pipe.2.html) [https://man7.org/linux/man-pages/man2/pipe.2.html] with the `O_CLOEXEC` flag set.

*New in version 1.41.0.*

# [uv\\_tty\\_t](#) — TTY handle

TTY handles represent a stream for the console.

[uv\\_tty\\_t](#) is a 'subclass' of [uv\\_stream\\_t](#).

## Data types

`uv_tty_t`

TTY handle type.

## Public members

N/A

## See also

The [uv\\_stream\\_t](#) members also apply.

## API

```
int uv_tty_init(uv\_loop\_t* loop, uv\_tty\_t* handle, uv\_file fd,
int unused)
```

Initialize a new TTY stream with the given file descriptor. Usually the file descriptor will be:

- 0 = stdin
- 1 = stdout
- 2 = stderr

On Unix this function will determine the path of the fd of the terminal using [ttyname\\_r\(3\)](#). [https://man7.org/linux/man-

pages/man3/ttyname\_r.3.html], open it, and use it if the passed file descriptor refers to a TTY. This lets libuv put the tty in non-blocking mode without affecting other processes that share the tty.

This function is not thread safe on systems that don't support ioctl TIOCGPTN or TIOCPTYGNAME, for instance OpenBSD and Solaris.

### **Note**

If reopening the TTY fails, libuv falls back to blocking writes.

*Changed in version 1.23.1::* the *readable* parameter is now unused and ignored. The correct value will now be auto-detected from the kernel.

*Changed in version 1.9.0::* the path of the TTY is determined by [ttyname\\_r\(3\)](https://man7.org/linux/man-pages/man3/ttyname_r.3.html) [https://man7.org/linux/man-pages/man3/ttyname\_r.3.html]. In earlier versions libuv opened `/dev/tty` instead.

*Changed in version 1.5.0::* trying to initialize a TTY stream with a file descriptor that refers to a file returns `UV_EINVAL` on UNIX.

int `uv_tty_set_mode`([uv\\_tty\\_t](#)\* *handle*, `uv_tty_mode_t` *mode*)

*Changed in version 1.2.0::* the mode is specified as a `uv_tty_mode_t` value.

Set the TTY using the specified terminal mode.

int `uv_tty_reset_mode`(void)

To be called when the program exits. Resets TTY settings to default values for the next process to take over.

This function is async signal-safe on Unix platforms but can fail with error code `UV_EBUSY` if you call it when execution is inside [uv\\_tty\\_set\\_mode\(\)](#).

```
int uv_tty_get_winsize(uv\_tty\_t* handle, int* width,  
int* height)
```

Gets the current Window size. On success it returns 0.

## See also

The [uv\\_stream\\_t](#) API functions also apply.

```
void uv_tty_set_vterm_state(uv_tty_vtermstate_t state)
```

Controls whether console virtual terminal sequences are processed by libuv or console. Useful in particular for enabling ConEmu support of ANSI X3.64 and Xterm 256 colors. Otherwise Windows10 consoles are usually detected automatically.

This function is only meaningful on Windows systems. On Unix it is silently ignored.

*New in version 1.33.0.*

```
int uv_tty_get_vterm_state(uv_tty_vtermstate_t* state)
```

Get the current state of whether console virtual terminal sequences are handled by libuv or the console.

This function is not implemented on Unix, where it returns `UV_ENOTSUP`.

*New in version 1.33.0.*

# uv\_udp\_t — UDP handle

UDP handles encapsulate UDP communication for both clients and servers.

## Data types

### uv\_udp\_t

UDP handle type.

### uv\_udp\_send\_t

UDP send request type.

### uv\_udp\_flags

Flags used in [uv\\_udp\\_bind\(\)](#) and [uv\\_udp\\_recv\\_cb..](#)

```
enum uv_udp_flags {
    /* Disables dual stack mode. */
    UV_UDP_IPV6ONLY = 1,
    /*
     * Indicates message was truncated because read buffer was
     * too small. The
     * remainder was discarded by the OS. Used in
     uv_udp_recv_cb.
     */
    UV_UDP_PARTIAL = 2,
    /*
     * Indicates if SO_REUSEADDR will be set when binding the
     handle in
     * uv_udp_bind.
     * This sets the SO_REUSEPORT socket flag on the BSDs and
     OS X. On other
     * Unix platforms, it sets the SO_REUSEADDR flag. What
     that means is that
     * multiple threads or processes can bind to the same
     address without error
     * (provided they all set the flag) but only the last one
     to bind will receive
     * any traffic, in effect "stealing" the port from the
```

```

previous listener.
    */
    UV_UDP_REUSEADDR = 4,
    /*
     * Indicates that the message was received by recvmsg,
    so the buffer provided
     * must not be freed by the recv_cb callback.
    */
    UV_UDP_MMSG_CHUNK = 8,
    /*
     * Indicates that the buffer provided has been fully
    utilized by recvmsg and
     * that it should now be freed by the recv_cb callback.
    When this flag is set
     * in uv_udp_recv_cb, nread will always be 0 and addr
    will always be NULL.
    */
    UV_UDP_MMSG_FREE = 16,
    /*
     * Indicates if IP_RECVERR/IPV6_RECVERR will be set when
    binding the handle.
     * This sets IP_RECVERR for IPv4 and IPV6_RECVERR for
    IPv6 UDP sockets on
     * Linux. This stops the Linux kernel from supressing
    some ICMP error messages
     * and enables full ICMP error reporting for faster
    failover.
     * This flag is no-op on platforms other than Linux.
    */
    UV_UDP_LINUX_RECVERR = 32,
    /*
     * Indicates that recvmsg should be used, if available.
    */
    UV_UDP_RECVMSG = 256
};

```

void (\*uv\_udp\_send\_cb)([uv\\_udp\\_send\\_t](#)\* req, int status)  
 Type definition for callback passed to [uv\\_udp\\_send\(\)](#), which is called after the data was sent.

void (\*uv\_udp\_recv\_cb)([uv\\_udp\\_t](#)\* handle, ssize\_t nread,  
 const [uv\\_buf\\_t](#)\* buf, const struct sockaddr\* addr,  
 unsigned flags)

Type definition for callback passed to [uv\\_udp\\_recv\\_start\(\)](#), which is called when the endpoint receives data.

- *handle*: UDP handle
- *nread*: Number of bytes that have been received. 0 if there is no more data to read. Note that 0 may also mean that an empty datagram was received (in this case *addr* is not NULL). < 0 if a transmission error was detected; if using [recvmsg\(2\)](#). [https://man7.org/linux/man-pages/man2/recvmsg.2.html] no more chunks will be received and the buffer can be freed safely.
- *buf*: [uv\\_buf\\_t](#) with the received data.
- *addr*: struct sockaddr\* containing the address of the sender. Can be NULL. Valid for the duration of the callback only.
- *flags*: One or more or'ed UV\_UDP\_\* constants.

The callee is responsible for freeing the buffer, libuv does not reuse it. The buffer may be a null buffer (where *buf->base* == NULL and *buf->len* == 0) on error.

When using [recvmsg\(2\)](#). [https://man7.org/linux/man-pages/man2/recvmsg.2.html], chunks will have the *UV\_UDP\_MMSG\_CHUNK* flag set, those must not be freed. If no errors occur, there will be a final callback with *nread* set to 0, *addr* set to NULL and the buffer pointing at the initially allocated data with the *UV\_UDP\_MMSG\_CHUNK* flag cleared and the *UV\_UDP\_MMSG\_FREE* flag set. If a UDP socket error occurs, *nread* will be < 0. In either scenario, the callee can now safely free the provided buffer.

*Changed in version 1.40.0:* added the *UV\_UDP\_MMSG\_FREE* flag.

## **Note**



The receive callback will be called with *nread* == 0 and *addr* == NULL when there is nothing to read, and with *nread* == 0 and *addr* != NULL when an empty UDP packet is received.

## Public members

`size_t uv_udp_t.send_queue_size`

Number of bytes queued for sending. This field strictly shows how much information is currently queued.

`size_t uv_udp_t.send_queue_count`

Number of send requests currently in the queue awaiting to be processed.

`uv_udp_t* uv_udp_send_t.handle`

UDP handle where this send request is taking place.

## See also

The [uv\\_handle\\_t](#) members also apply.

## API

`int uv_udp_init(uv_loop_t* loop, uv_udp_t* handle)`

Initialize a new UDP handle. The actual socket is created lazily. Returns 0 on success.

`int uv_udp_init_ex(uv_loop_t* loop, uv_udp_t* handle, unsigned int flags)`

Initialize the handle with the specified flags. The lower 8 bits of the *flags* parameter are used as the socket domain. A socket will be created for the given domain. If the

specified domain is `AF_UNSPEC` no socket is created, just like [uv\\_udp\\_init\(\)](#).

The remaining bits can be used to set one of these flags:

- `UV_UDP_RECVMMSG`: if set, and the platform supports it, [recvmsg\(2\)](#) [<https://man7.org/linux/man-pages/man2/recvmsg.2.html>] will be used.

*New in version 1.7.0.*

*Changed in version 1.37.0:* added the `UV_UDP_RECVMMSG` flag.

`int uv_udp_open(uv\_udp\_t\* handle, uv\_os\_sock\_t sock)`

Opens an existing file descriptor or Windows SOCKET as a UDP handle.

Unix only: The only requirement of the `sock` argument is that it follows the datagram contract (works in unconnected mode, supports `sendmsg()/recvmsg()`, etc). In other words, other datagram-type sockets like raw sockets or netlink sockets can also be passed to this function.

*Changed in version 1.2.1:* the file descriptor is set to non-blocking mode.

### **Note**

The passed file descriptor or SOCKET is not checked for its type, but it's required that it represents a valid datagram socket.

`int uv_udp_bind(uv\_udp\_t\* handle, const struct sockaddr* addr, unsigned int flags)`

Bind the UDP handle to an IP address and port.

- Parameters:**
- **handle** - UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#).
  - **addr** - *struct sockaddr\_in* or *struct sockaddr\_in6* with the address and port to bind to.
  - **flags** - Indicate how the socket will be bound, `UV_UDP_IPV6ONLY`, `UV_UDP_REUSEADDR`, and `UV_UDP_RECVERR` are supported.
- Returns:** 0 on success, or an error code < 0 on failure.

```
int uv_udp_connect(uv\_udp\_t\* handle, const struct  
sockaddr* addr)
```

Associate the UDP handle to a remote address and port, so every message sent by this handle is automatically sent to that destination. Calling this function with a *NULL* *addr* disconnects the handle. Trying to call *uv\_udp\_connect()* on an already connected handle will result in an *UV\_EISCONN* error. Trying to disconnect a handle that is not connected will return an *UV\_ENOTCONN* error.

- Parameters:**
- **handle** - UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#).
  - **addr** - *struct sockaddr\_in* or *struct sockaddr\_in6* with the address and port to associate to.
- Returns:** 0 on success, or an error code < 0 on failure.

*New in version 1.27.0.*

```
int uv_udp_getpeername(const uv\_udp\_t\* handle, struct  
sockaddr* name, int* namelen)
```

Get the remote IP and port of the UDP handle on connected UDP handles. On unconnected handles, it returns *UV\_ENOTCONN*.

- Parameters:**
- **handle** – UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#) and bound.
  - **name** – Pointer to the structure to be filled with the address data. In order to support IPv4 and IPv6 *struct sockaddr\_storage* should be used.
  - **namelen** – On input it indicates the dat of the *name* field. On output it indicates how much of it was filled.

**Returns:** 0 on success, or an error code < 0 on failure

*New in version 1.27.0.*

```
int uv_udp_getsockname(const uv_udp_t* handle, struct  
sockaddr* name, int* namelen)
```

Get the local IP and port of the UDP handle.

- Parameters:**
- **handle** – UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#) and bound.
  - **name** – Pointer to the structure to be filled with the address data. In order to support IPv4 and IPv6 *struct sockaddr\_storage* should be used.
  - **namelen** – On input it indicates the dat of the *name* field. On output it indicates how much of it was filled.

**Returns:** 0 on success, or an error code < 0 on failure.

```
int uv_udp_set_membership(uv_udp_t* handle, const  
char* multicast_addr, const char* interface_addr,  
uv_membership membership)
```

Set membership for a multicast address

- Parameters:**
- **handle** - UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#).
  - **multicast\_addr** - Multicast address to set membership for.
  - **interface\_addr** - Interface address.
  - **membership** - Should be UV\_JOIN\_GROUP or UV\_LEAVE\_GROUP.
- Returns:** 0 on success, or an error code < 0 on failure.

```
int uv_udp_set_source_membership(uv_udp_t* handle, const char* multicast_addr, const char* interface_addr, const char* source_addr, uv_membership membership)
```

Set membership for a source-specific multicast group.

- Parameters:**
- **handle** - UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#).
  - **multicast\_addr** - Multicast address to set membership for.
  - **interface\_addr** - Interface address.
  - **source\_addr** - Source address.
  - **membership** - Should be UV\_JOIN\_GROUP or UV\_LEAVE\_GROUP.
- Returns:** 0 on success, or an error code < 0 on failure.

*New in version 1.32.0.*

```
int uv_udp_set_multicast_loop(uv_udp_t* handle, int on)
```

Set IP multicast loop flag. Makes multicast packets loop back to local sockets.

- Parameters:**
- **handle** - UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#).
  - **on** - 1 for on, 0 for off.
- Returns:** 0 on success, or an error code < 0 on failure.

int `uv_udp_set_multicast_ttl`([uv\\_udp\\_t\\*](#) *handle*, int *ttl*)

Set the multicast ttl.

**Parameters:**

- **handle** - UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#).

- **ttl** - 1 through 255.

**Returns:** 0 on success, or an error code < 0 on failure.

int `uv_udp_set_multicast_interface`([uv\\_udp\\_t\\*](#) *handle*, const char\* *interface\_addr*)

Set the multicast interface to send or receive data on.

**Parameters:**

- **handle** - UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#).

- **interface\_addr** - interface address.

**Returns:** 0 on success, or an error code < 0 on failure.

int `uv_udp_set_broadcast`([uv\\_udp\\_t\\*](#) *handle*, int *on*)

Set broadcast on or off.

**Parameters:**

- **handle** - UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#).

- **on** - 1 for on, 0 for off.

**Returns:** 0 on success, or an error code < 0 on failure.

int `uv_udp_set_ttl`([uv\\_udp\\_t\\*](#) *handle*, int *ttl*)

Set the time to live.

**Parameters:**

- **handle** - UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#).

- **ttl** - 1 through 255.

**Returns:** 0 on success, or an error code < 0 on failure.

```
int uv_udp_send(uv\_udp\_send\_t* req, uv\_udp\_t* handle, const
uv\_buf\_t bufs[], unsigned int nbufs, const struct
sockaddr* addr, uv\_udp\_send\_cb send_cb)
```

Send data over the UDP socket. If the socket has not previously been bound with [uv\\_udp\\_bind\(\)](#) it will be bound to 0.0.0.0 (the “all interfaces” IPv4 address) and a random port number.

On Windows if the *addr* is initialized to point to an unspecified address (0.0.0.0 or ::) it will be changed to point to localhost. This is done to match the behavior of Linux systems.

For connected UDP handles, *addr* must be set to *NULL*, otherwise it will return *UV\_EISCONN* error.

For connectionless UDP handles, *addr* cannot be *NULL*, otherwise it will return *UV\_EDESTADDRREQ* error.

- Parameters:**
- **req** – UDP request handle. Need not be initialized.
  - **handle** – UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#).
  - **bufs** – List of buffers to send.
  - **nbufs** – Number of buffers in *bufs*.
  - **addr** – *struct sockaddr\_in* or *struct sockaddr\_in6* with the address and port of the remote peer.
  - **send\_cb** – Callback to invoke when the data has been sent out.
- Returns:** 0 on success, or an error code < 0 on failure.

*Changed in version 1.19.0:* added 0.0.0.0 and :: to localhost mapping

*Changed in version 1.27.0:* added support for connected sockets

```
int uv_udp_try_send(uv\_udp\_t* handle, const uv\_buf\_t bufs[],  
unsigned int nbufs, const struct sockaddr* addr)
```

Same as [uv\\_udp\\_send\(\)](#), but won't queue a send request if it can't be completed immediately.

For connected UDP handles, *addr* must be set to *NULL*, otherwise it will return *UV\_EISCONN* error.

For connectionless UDP handles, *addr* cannot be *NULL*, otherwise it will return *UV\_EDESTADDRREQ* error.

**Returns:**             $\geq 0$ : number of bytes sent (it matches the given buffer size).  $< 0$ : negative error code (*UV\_EAGAIN* is returned when the message can't be sent immediately).

*Changed in version 1.27.0:* added support for connected sockets

```
int uv_udp_recv_start(uv\_udp\_t* handle, uv\_alloc\_cb alloc_cb,  
uv\_udp\_recv\_cb recv_cb)
```

Prepare for receiving data. If the socket has not previously been bound with [uv\\_udp\\_bind\(\)](#) it is bound to 0.0.0.0 (the "all interfaces" IPv4 address) and a random port number.

**Parameters:**

- **handle** - UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#).
- **alloc\_cb** - Callback to invoke when temporary storage is needed.
- **recv\_cb** - Callback to invoke with received data.

**Returns:**            0 on success, or an error code  $< 0$  on failure.



## Note

When using [recvmsg\(2\)](https://man7.org/linux/man-pages/man2/recvmsg.2.html), the number of messages received at a time is limited by the number of max size dgrams that will fit into the buffer allocated in *alloc\_cb*, and *suggested\_size* in *alloc\_cb* for *udp\_recv* is always set to the size of 1 max size dgram.

*Changed in version 1.35.0:* added support for [recvmsg\(2\)](https://man7.org/linux/man-pages/man2/recvmsg.2.html) on supported platforms). The use of this feature requires a buffer larger than 2 \* 64KB to be passed to *alloc\_cb*.

*Changed in version 1.37.0:* [recvmsg\(2\)](https://man7.org/linux/man-pages/man2/recvmsg.2.html) support is no longer enabled implicitly, it must be explicitly requested by passing the *UV\_UDP\_RECVMSG* flag to [uv\\_udp\\_init\\_ex\(\)](#).

*Changed in version 1.39.0:* [uv\\_udp\\_using\\_recvmsg\(\)](#) can be used in *alloc\_cb* to determine if a buffer sized for use with [recvmsg\(2\)](https://man7.org/linux/man-pages/man2/recvmsg.2.html) should be allocated for the current handle/platform.

`int uv_udp_using_recvmsg(uv_udp_t* handle)`

Returns 1 if the UDP handle was created with the *UV\_UDP\_RECVMSG* flag and the platform supports [recvmsg\(2\)](https://man7.org/linux/man-pages/man2/recvmsg.2.html), 0 otherwise.

*New in version 1.39.0.*

`int uv_udp_recv_stop(uv_udp_t* handle)`

Stop listening for incoming datagrams.

**Parameters:**     • **handle** – UDP handle. Should have been initialized with [uv\\_udp\\_init\(\)](#).

**Returns:**         0 on success, or an error code < 0 on failure.

size\_t [uv\\_udp\\_get\\_send\\_queue\\_size](#)(const [uv\\_udp\\_t](#)\* handle)  
Returns *handle->send\_queue\_size*.

*New in version 1.19.0.*

size\_t [uv\\_udp\\_get\\_send\\_queue\\_count](#)(const [uv\\_udp\\_t](#)\* handle)  
Returns *handle->send\_queue\_count*.

*New in version 1.19.0.*

## See also

The [uv\\_handle\\_t](#) API functions also apply.

# uv\_fs\_event\_t — FS Event handle

FS Event handles allow the user to monitor a given path for changes, for example, if the file was renamed or there was a generic change in it. This handle uses the best backend for the job on each platform.

## Note

For AIX, the non default IBM bos.ahafs package has to be installed. The AIX Event Infrastructure file system (ahafs) has some limitations:

- ahafs tracks monitoring per process and is not thread safe. A separate process must be spawned for each monitor for the same event.
- Events for file modification (writing to a file) are not received if only the containing folder is watched.

See [documentation](https://developer.ibm.com/articles/au-aix_event_infrastructure/) [https://developer.ibm.com/articles/au-aix\_event\_infrastructure/] for more details.

The z/OS file system events monitoring infrastructure does not notify of file creation/deletion within a directory that is being monitored. See the [IBM Knowledge centre](https://www.ibm.com/support/knowledgecenter/en/SSLTBW_2.2.0/com.ibm.zos.v2r1.bpxb100/ioc.htm) [https://www.ibm.com/support/knowledgecenter/en/SSLTBW\_2.2.0/com.ibm.zos.v2r1.bpxb100/ioc.htm] for more details.

## Data types

## uv\_fs\_event\_t

FS Event handle type.

void (\*uv\_fs\_event\_cb)(uv\_fs\_event\_t\* handle, const char\* filename, int events, int status)

Callback passed to [uv\\_fs\\_event\\_start\(\)](#) which will be called repeatedly after the handle is started. If the handle was started with a directory the *filename* parameter will be a relative path to a file contained in the directory. The *events* parameter is an ORed mask of [uv\\_fs\\_event](#) elements.

## uv\_fs\_event

Event types that [uv\\_fs\\_event\\_t](#) handles monitor.

```
enum uv_fs_event {
    UV_RENAME = 1,
    UV_CHANGE = 2
};
```

## uv\_fs\_event\_flags

Flags that can be passed to [uv\\_fs\\_event\\_start\(\)](#) to control its behavior.

```
enum uv_fs_event_flags {
    /*
     * By default, if the fs event watcher is given a
     * directory name, we will
     * watch for all events in that directory. This flags
     * overrides this behavior
     * and makes fs_event report only changes to the
     * directory entry itself. This
     * flag does not affect individual files watched.
     * This flag is currently not implemented yet on any
     * backend.
     */
    UV_FS_EVENT_WATCH_ENTRY = 1,
    /*
     * By default uv_fs_event will try to use a kernel
     * interface such as inotify
```

```

    * or kqueue to detect events. This may not work on
remote file systems such
    * as NFS mounts. This flag makes fs_event fall back to
calling stat() on a
    * regular interval.
    * This flag is currently not implemented yet on any
backend.
    */
    UV_FS_EVENT_STAT = 2,
    /*
    * By default, event watcher, when watching directory, is
not registering
    * (is ignoring) changes in its subdirectories.
    * This flag will override this behaviour on platforms
that support it.
    */
    UV_FS_EVENT_RECURSIVE = 4
};

```

## Public members

N/A

## See also

The [uv\\_handle\\_t](#) members also apply.

## API

```
int uv_fs_event_init(uv\_loop\_t* loop,
uv\_fs\_event\_t* handle)
```

Initialize the handle.

```
int uv_fs_event_start(uv\_fs\_event\_t* handle,
uv\_fs\_event\_cb cb, const char* path, unsigned int flags)
```

Start the handle with the given callback, which will watch the specified *path* for changes. *flags* can be an ORed

mask of [uv\\_fs\\_event\\_flags](#).

### **Note**

Currently the only supported flag is `UV_FS_EVENT_RECURSIVE` and only on OSX and Windows.

`int uv_fs_event_stop(uv\_fs\_event\_t* handle)`

Stop the handle, the callback will no longer be called.

`int uv_fs_event_getpath(uv\_fs\_event\_t* handle, char* buffer, size_t* size)`

Get the path being monitored by the handle. The buffer must be preallocated by the user. Returns 0 on success or an error code  $< 0$  in case of failure. On success, *buffer* will contain the path and *size* its length. If the buffer is not big enough `UV_ENOBUFS` will be returned and *size* will be set to the required size, including the null terminator.

*Changed in version 1.3.0:* the returned length no longer includes the terminating null byte, and the buffer is not null terminated.

*Changed in version 1.9.0:* the returned length includes the terminating null byte on `UV_ENOBUFS`, and the buffer is null terminated on success.

### **See also**

The [uv\\_handle\\_t](#) API functions also apply.

# [uv\\_fs\\_poll\\_t](#) — FS Poll handle

FS Poll handles allow the user to monitor a given path for changes. Unlike [uv\\_fs\\_event\\_t](#), fs poll handles use *stat* to detect when a file has changed so they can work on file systems where fs event handles can't.

## Data types

[uv\\_fs\\_poll\\_t](#)

FS Poll handle type.

`void (*uv\_fs\_poll\_cb)(uv\_fs\_poll\_t* handle, int status, const uv\_stat\_t* prev, const uv\_stat\_t* curr)`

Callback passed to [uv\\_fs\\_poll\\_start\(\)](#) which will be called repeatedly after the handle is started, when any change happens to the monitored path.

The callback is invoked with *status* < 0 if *path* does not exist or is inaccessible. The watcher is *not* stopped but your callback is not called again until something changes (e.g. when the file is created or the error reason changes).

When *status* == 0, the callback receives pointers to the old and new [uv\\_stat\\_t](#) structs. They are valid for the duration of the callback only.

## Public members

N/A

## See also

The [uv\\_handle\\_t](#) members also apply.

## API

int [uv\\_fs\\_poll\\_init](#)([uv\\_loop\\_t](#)\* *loop*, [uv\\_fs\\_poll\\_t](#)\* *handle*)  
Initialize the handle.

int [uv\\_fs\\_poll\\_start](#)([uv\\_fs\\_poll\\_t](#)\* *handle*,  
[uv\\_fs\\_poll\\_cb](#) *poll\_cb*, const char\* *path*, unsigned  
int *interval*)

Check the file at *path* for changes every *interval* milliseconds.

### Note

For maximum portability, use multi-second intervals. Sub-second intervals will not detect all changes on many file systems.

int [uv\\_fs\\_poll\\_stop](#)([uv\\_fs\\_poll\\_t](#)\* *handle*)  
Stop the handle, the callback will no longer be called.

int [uv\\_fs\\_poll\\_getpath](#)([uv\\_fs\\_poll\\_t](#)\* *handle*, char\* *buffer*,  
size\_t\* *size*)

Get the path being monitored by the handle. The buffer must be preallocated by the user. Returns 0 on success or an error code < 0 in case of failure. On success, *buffer* will contain the path and *size* its length. If the buffer is not big enough *UV\_ENOBUFS* will be returned and *size* will be set to the required size.



*Changed in version 1.3.0:* the returned length no longer includes the terminating null byte, and the buffer is not null terminated.

*Changed in version 1.9.0:* the returned length includes the terminating null byte on `UV_ENOBUFS`, and the buffer is null terminated on success.

## **See also**

The [uv\\_handle\\_t](#) API functions also apply.

# File system operations

libuv provides a wide variety of cross-platform sync and async file system operations. All functions defined in this document take a callback, which is allowed to be NULL. If the callback is NULL the request is completed synchronously, otherwise it will be performed asynchronously.

All file operations are run on the threadpool. See [Thread pool work scheduling](#) for information on the threadpool size.

## Note

On Windows `uv_fs_*` functions use utf-8 encoding.

## Data types

### `uv_fs_t`

File system request type.

### `uv_timespec_t`

Portable equivalent of `struct timespec`.

```
typedef struct {
    long tv_sec;
    long tv_nsec;
} uv_timespec_t;
```

### `uv_stat_t`

Portable equivalent of `struct stat`.

```
typedef struct {
    uint64_t st_dev;
```

```

uint64_t st_mode;
uint64_t st_nlink;
uint64_t st_uid;
uint64_t st_gid;
uint64_t st_rdev;
uint64_t st_ino;
uint64_t st_size;
uint64_t st_blksize;
uint64_t st_blocks;
uint64_t st_flags;
uint64_t st_gen;
uv_timespec_t st_atim;
uv_timespec_t st_mtim;
uv_timespec_t st_ctim;
uv_timespec_t st_birthtim;
} uv_stat_t;

```

### **uv\_statfs\_t**

Reduced cross platform equivalent of struct statfs. Used in [uv\\_fs\\_statfs\(\)](#).

```

typedef struct uv_statfs_s {
    uint64_t f_type;
    uint64_t f_bsize;
    uint64_t f_blocks;
    uint64_t f_bfree;
    uint64_t f_bavail;
    uint64_t f_files;
    uint64_t f_ffree;
    uint64_t f_spare[4];
} uv_statfs_t;

```

### **uv\_dir\_t**

Data type used for streaming directory iteration. Used by [uv\\_fs\\_opendir\(\)](#), [uv\\_fs\\_readdir\(\)](#), and [uv\\_fs\\_closedir\(\)](#).

*dirents* represents a user provided array of *uv\_dirent\_t*'s used to hold results. *nentries* is the user provided maximum array size of *dirents*.

```

typedef struct uv_dir_s {
    uv_dirent_t* dirents;

```

```
    size_t nentries;  
} uv_dir_t;
```

## Public members

[uv\\_loop\\_t](#)\* **uv\_fs\_t.loop**

Loop that started this request and where completion will be reported. Readonly.

uv\_fs\_type **uv\_fs\_t.fs\_type**

FS request type.

const char\* **uv\_fs\_t.path**

Path affecting the request.

ssize\_t **uv\_fs\_t.result**

Result of the request. < 0 means error, success otherwise. On requests such as [uv\\_fs\\_read\(\)](#) or [uv\\_fs\\_write\(\)](#) it indicates the amount of data that was read or written, respectively.

[uv\\_stat\\_t](#) **uv\_fs\_t.statbuf**

Stores the result of [uv\\_fs\\_stat\(\)](#) and other stat requests.

void\* **uv\_fs\_t.ptr**

Stores the result of [uv\\_fs\\_readlink\(\)](#) and [uv\\_fs\\_realpath\(\)](#) and serves as an alias to *statbuf*.

## See also

The [uv\\_req\\_t](#) members also apply.

## API

void `uv_fs_req_cleanup`([uv\\_fs\\_t\\*](#) req)

Cleanup request. Must be called after a request is finished to deallocate any memory libuv might have allocated.

int `uv_fs_close`([uv\\_loop\\_t\\*](#) loop, [uv\\_fs\\_t\\*](#) req, [uv\\_file](#) file, [uv\\_fs\\_cb](#) cb)

Equivalent to [close\(2\)](#). [<https://man7.org/linux/man-pages/man2/close.2.html>].

int `uv_fs_open`([uv\\_loop\\_t\\*](#) loop, [uv\\_fs\\_t\\*](#) req, const char\* path, int flags, int mode, [uv\\_fs\\_cb](#) cb)

Equivalent to [open\(2\)](#). [<https://man7.org/linux/man-pages/man2/open.2.html>].

### Note

On Windows libuv uses *CreateFileW* and thus the file is always opened in binary mode. Because of this the `O_BINARY` and `O_TEXT` flags are not supported.

int `uv_fs_read`([uv\\_loop\\_t\\*](#) loop, [uv\\_fs\\_t\\*](#) req, [uv\\_file](#) file, const [uv\\_buf\\_t](#) bufs[], unsigned int nbufs, int64\_t offset, [uv\\_fs\\_cb](#) cb)

Equivalent to [preadv\(2\)](#). [<https://man7.org/linux/man-pages/man2/preadv.2.html>].

### Warning

On Windows, under non-MSVC environments (e.g. when GCC or Clang is used to build libuv), files opened using `UV_FS_O_FILEMAP` may cause a fatal crash if the memory mapped read operation fails.

```
int uv_fs_unlink(uv\_loop\_t* loop, uv\_fs\_t* req, const
char* path, uv_fs_cb cb)
```

Equivalent to [unlink\(2\)](https://man7.org/linux/man-pages/man2/unlink.2.html). [https://man7.org/linux/man-pages/man2/unlink.2.html].

```
int uv_fs_write(uv\_loop\_t* loop, uv\_fs\_t* req, uv\_file file,
const uv\_buf\_t bufs[], unsigned int nbufs, int64_t offset,
uv_fs_cb cb)
```

Equivalent to [pwritev\(2\)](https://man7.org/linux/man-pages/man2/pwritev.2.html). [https://man7.org/linux/man-pages/man2/pwritev.2.html].

### Warning

On Windows, under non-MSVC environments (e.g. when GCC or Clang is used to build libuv), files opened using UV\_FS\_O\_FILEMAP may cause a fatal crash if the memory mapped write operation fails.

```
int uv_fs_mkdir(uv\_loop\_t* loop, uv\_fs\_t* req, const
char* path, int mode, uv_fs_cb cb)
```

Equivalent to [mkdir\(2\)](https://man7.org/linux/man-pages/man2/mkdir.2.html). [https://man7.org/linux/man-pages/man2/mkdir.2.html].

### Note

*mode* is currently not implemented on Windows.

```
int uv_fs_mkdtemp(uv\_loop\_t* loop, uv\_fs\_t* req, const
char* tpl, uv_fs_cb cb)
```

Equivalent to [mkdtemp\(3\)](https://man7.org/linux/man-pages/man3/mkdtemp.3.html). [https://man7.org/linux/man-pages/man3/mkdtemp.3.html]. The result can be found as a null terminated string at *req->path*.

```
int uv_fs_mkstemp(uv\_loop\_t* loop, uv\_fs\_t* req, const
char* tpl, uv_fs_cb cb)
```

Equivalent to [mkstemp\(3\)](#) [<https://man7.org/linux/man-pages/man3/mkstemp.3.html>]. The created file path can be found as a null terminated string at `req->path`. The file descriptor can be found as an integer at `req->result`.

*New in version 1.34.0.*

```
int uv_fs_rmdir(uv\_loop\_t* loop, uv\_fs\_t* req, const
char* path, uv_fs_cb cb)
```

Equivalent to [rmdir\(2\)](#) [<https://man7.org/linux/man-pages/man2/rmdir.2.html>].

```
int uv_fs_opendir(uv\_loop\_t* loop, uv\_fs\_t* req, const
char* path, uv_fs_cb cb)
```

Opens `path` as a directory stream. On success, a `uv_dir_t` is allocated and returned via `req->ptr`. This memory is not freed by `uv_fs_req_cleanup()`, although `req->ptr` is set to `NULL`. The allocated memory must be freed by calling `uv_fs_closedir()`. On failure, no memory is allocated.

The contents of the directory can be iterated over by passing the resulting `uv_dir_t` to `uv_fs_readdir()`.

*New in version 1.28.0.*

```
int uv_fs_closedir(uv\_loop\_t* loop, uv\_fs\_t* req,
uv\_dir\_t* dir, uv_fs_cb cb)
```

Closes the directory stream represented by `dir` and frees the memory allocated by `uv_fs_opendir()`.

*New in version 1.28.0.*

```
int uv_fs_readdir(uv\_loop\_t* loop, uv\_fs\_t* req, uv\_dir\_t* dir,
uv_fs_cb cb)
```

Iterates over the directory stream, *dir*, returned by a successful *uv\_fs\_opendir()* call. Prior to invoking *uv\_fs\_readdir()*, the caller must set *dir->dirents* and *dir->nentries*, representing the array of [uv\\_dirent\\_t](#) elements used to hold the read directory entries and its size.

On success, the result is an integer  $\geq 0$  representing the number of entries read from the stream.

*New in version 1.28.0.*

### **Warning**

*uv\_fs\_readdir()* is not thread safe.

### **Note**

This function does not return the “.” and “..” entries.

### **Note**

On success this function allocates memory that must be freed using *uv\_fs\_req\_cleanup()*. *uv\_fs\_req\_cleanup()* must be called before closing the directory with *uv\_fs\_closedir()*.

```
int uv_fs_scandir(uv\_loop\_t* loop, uv\_fs\_t* req, const
char* path, int flags, uv_fs_cb cb)
```

```
int uv_fs_scandir_next(uv\_fs\_t* req, uv\_dirent\_t* ent)
```

Equivalent to [scandir\(3\)](#) [<https://man7.org/linux/man-pages/man3/scandir.3.html>], with a slightly different API. Once the callback for the request is called, the user can use



[uv\\_fs\\_scandir\\_next\(\)](#) to get *ent* populated with the next directory entry data. When there are no more entries UV\_EOF will be returned.

### Note

Unlike *scandir(3)*, this function does not return the “.” and “..” entries.

### Note

On Linux, getting the type of an entry is only supported by some file systems (btrfs, ext2, ext3 and ext4 at the time of this writing), check the [getdents\(2\)](#).

[<https://man7.org/linux/man-pages/man2/getdents.2.html>] man page.

```
int uv_fs_stat(uv\_loop\_t\* loop, uv\_fs\_t\* req, const char* path, uv_fs_cb cb)
```

```
int uv_fs_fstat(uv\_loop\_t\* loop, uv\_fs\_t\* req, uv\_file file, uv_fs_cb cb)
```

```
int uv_fs_lstat(uv\_loop\_t\* loop, uv\_fs\_t\* req, const char* path, uv_fs_cb cb)
```

Equivalent to [stat\(2\)](#) [<https://man7.org/linux/man-pages/man2/stat.2.html>], [fstat\(2\)](#) [<https://man7.org/linux/man-pages/man2/fstat.2.html>] and [lstat\(2\)](#) [<https://man7.org/linux/man-pages/man2/lstat.2.html>] respectively.

```
int uv_fs_statfs(uv\_loop\_t\* loop, uv\_fs\_t\* req, const char* path, uv_fs_cb cb)
```

Equivalent to [statfs\(2\)](#) [<https://man7.org/linux/man-pages/man2/statfs.2.html>]. On success, a *uv\_statfs\_t* is

allocated and returned via `req->ptr`. This memory is freed by `uv_fs_req_cleanup()`.

### **Note**

Any fields in the resulting `uv_statfs_t` that are not supported by the underlying operating system are set to zero.

*New in version 1.31.0.*

```
int uv_fs_rename(uv\_loop\_t\* loop, uv\_fs\_t\* req, const char* path, const char* new_path, uv_fs_cb cb)
```

Equivalent to [rename\(2\)](#). [<https://man7.org/linux/man-pages/man2/rename.2.html>].

```
int uv_fs_fsync(uv\_loop\_t\* loop, uv\_fs\_t\* req, uv\_file file, uv_fs_cb cb)
```

Equivalent to [fsync\(2\)](#). [<https://man7.org/linux/man-pages/man2/fsync.2.html>].

### **Note**

For AIX, `uv_fs_fsync` returns `UV_EBADF` on file descriptors referencing non regular files.

```
int uv_fs_fdatasync(uv\_loop\_t\* loop, uv\_fs\_t\* req, uv\_file file, uv_fs_cb cb)
```

Equivalent to [fdatasync\(2\)](#). [<https://man7.org/linux/man-pages/man2/fdatasync.2.html>].

```
int uv_fs_ftruncate(uv\_loop\_t\* loop, uv\_fs\_t\* req, uv\_file file, int64_t offset, uv_fs_cb cb)
```

Equivalent to [ftruncate\(2\)](https://man7.org/linux/man-pages/man2/ftruncate.2.html). [https://man7.org/linux/man-pages/man2/ftruncate.2.html].

```
int uv_fs_copyfile(uv_loop_t* loop, uv_fs_t* req, const char* path, const char* new_path, int flags, uv_fs_cb cb)
```

Copies a file from *path* to *new\_path*. Supported *flags* are described below.

- *UV\_FS\_COPYFILE\_EXCL*: If present, *uv\_fs\_copyfile()* will fail with *UV\_EEXIST* if the destination path already exists. The default behavior is to overwrite the destination if it exists.
- *UV\_FS\_COPYFILE\_FICLONE*: If present, *uv\_fs\_copyfile()* will attempt to create a copy-on-write reflink. If the underlying platform does not support copy-on-write, or an error occurs while attempting to use copy-on-write, a fallback copy mechanism based on [uv\\_fs\\_sendfile\(\)](#) is used.
- *UV\_FS\_COPYFILE\_FICLONE\_FORCE*: If present, *uv\_fs\_copyfile()* will attempt to create a copy-on-write reflink. If the underlying platform does not support copy-on-write, or an error occurs while attempting to use copy-on-write, then an error is returned.

### **Warning**

If the destination path is created, but an error occurs while copying the data, then the destination path is removed. There is a brief window of time between closing and removing the file where another process could access the file.

*New in version 1.14.0.*

*Changed in version 1.20.0: *UV\_FS\_COPYFILE\_FICLONE* and *UV\_FS\_COPYFILE\_FICLONE\_FORCE* are supported.*

*Changed in version 1.33.0:* If an error occurs while using `UV_FS_COPYFILE_FICLONE_FORCE`, that error is returned. Previously, all errors were mapped to `UV_ENOTSUP`.

```
int uv_fs_sendfile(uv\_loop\_t\* loop, uv\_fs\_t\* req,  
uv\_file out_fd, uv\_file in_fd, int64_t in_offset, size_t length,  
uv_fs_cb cb)
```

Limited equivalent to [sendfile\(2\)](#). [<https://man7.org/linux/man-pages/man2/sendfile.2.html>].

```
int uv_fs_access(uv\_loop\_t\* loop, uv\_fs\_t\* req, const  
char* path, int mode, uv_fs_cb cb)
```

Equivalent to [access\(2\)](#). [<https://man7.org/linux/man-pages/man2/access.2.html>] on Unix. Windows uses `GetFileAttributesW()`.

```
int uv_fs_chmod(uv\_loop\_t\* loop, uv\_fs\_t\* req, const  
char* path, int mode, uv_fs_cb cb)
```

```
int uv_fs_fchmod(uv\_loop\_t\* loop, uv\_fs\_t\* req, uv\_file file,  
int mode, uv_fs_cb cb)
```

Equivalent to [chmod\(2\)](#). [<https://man7.org/linux/man-pages/man2/chmod.2.html>] and [fchmod\(2\)](#). [<https://man7.org/linux/man-pages/man2/fchmod.2.html>] respectively.

```
int uv_fs_utime(uv\_loop\_t\* loop, uv\_fs\_t\* req, const  
char* path, double atime, double mtime, uv_fs_cb cb)
```

```
int uv_fs_futime(uv\_loop\_t\* loop, uv\_fs\_t\* req, uv\_file file,  
double atime, double mtime, uv_fs_cb cb)
```

```
int uv_fs_lutime(uv\_loop\_t\* loop, uv\_fs\_t\* req, const  
char* path, double atime, double mtime, uv_fs_cb cb)
```

Equivalent to [utime\(2\)](https://man7.org/linux/man-pages/man2/utime.2.html) [https://man7.org/linux/man-pages/man2/utime.2.html], [futimes\(3\)](https://man7.org/linux/man-pages/man3/futimes.3.html) [https://man7.org/linux/man-pages/man3/futimes.3.html] and [lutimes\(3\)](https://man7.org/linux/man-pages/man3/lutimes.3.html) [https://man7.org/linux/man-pages/man3/lutimes.3.html] respectively.

### Note

z/OS: `uv_fs_lutime()` is not implemented for z/OS. It can still be called but will return `UV_ENOSYS`.

### Note

AIX: `uv_fs_futime()` and `uv_fs_lutime()` functions only work for AIX 7.1 and newer. They can still be called on older versions but will return `UV_ENOSYS`.

*Changed in version 1.10.0:* sub-second precision is supported on Windows

```
int uv_fs_link(uv_loop_t* loop, uv_fs_t* req, const char* path, const char* new_path, uv_fs_cb cb)
```

Equivalent to [link\(2\)](https://man7.org/linux/man-pages/man2/link.2.html) [https://man7.org/linux/man-pages/man2/link.2.html].

```
int uv_fs_symlink(uv_loop_t* loop, uv_fs_t* req, const char* path, const char* new_path, int flags, uv_fs_cb cb)
```

Equivalent to [symlink\(2\)](https://man7.org/linux/man-pages/man2/symlink.2.html) [https://man7.org/linux/man-pages/man2/symlink.2.html].

### Note

On Windows the `flags` parameter can be specified to control how the symlink will be created:

- `UV_FS_SYMLINK_DIR`: indicates that `path` points to a directory.

- `UV_FS_SYMLINK_JUNCTION`: request that the symlink is created using junction points.

```
int uv_fs_readlink(uv\_loop\_t\* loop, uv\_fs\_t\* req, const char* path, uv_fs_cb cb)
```

Equivalent to [readlink\(2\)](#) [<https://man7.org/linux/man-pages/man2/readlink.2.html>]. The resulting string is stored in `req->ptr`.

```
int uv_fs_realpath(uv\_loop\_t\* loop, uv\_fs\_t\* req, const char* path, uv_fs_cb cb)
```

Equivalent to [realpath\(3\)](#) [<https://man7.org/linux/man-pages/man3/realpath.3.html>] on Unix. Windows uses [GetFinalPathNameByHandle](#) [<https://docs.microsoft.com/en-us/windows/win32/api/fileapi/nf-fileapi-getfinalpathnamebyhandlea>]. The resulting string is stored in `req->ptr`.

## Warning

This function has certain platform-specific caveats that were discovered when used in Node.

- macOS and other BSDs: this function will fail with `UV_ELOOP` if more than 32 symlinks are found while resolving the given path. This limit is hardcoded and cannot be sidestepped.
- Windows: while this function works in the common case, there are a number of corner cases where it doesn't:
  - Paths in ramdisk volumes created by tools which sidestep the Volume Manager (such as `ImDisk`) cannot be resolved.
  - Inconsistent casing when using drive letters.
  - Resolved path bypasses subst'd drives.

While this function can still be used, it's not recommended if scenarios such as the above need to be supported.

The background story and some more details on these issues can be checked [here](#)

[<https://github.com/nodejs/node/issues/7726>].

## Note

This function is not implemented on Windows XP and Windows Server 2003. On these systems, UV\_ENOSYS is returned.

*New in version 1.8.0.*

```
int uv_fs_chown(uv\_loop\_t\* loop, uv\_fs\_t\* req, const char* path, uv_uid_t uid, uv_gid_t gid, uv_fs_cb cb)
```

```
int uv_fs_fchown(uv\_loop\_t\* loop, uv\_fs\_t\* req, uv\_file file, uv_uid_t uid, uv_gid_t gid, uv_fs_cb cb)
```

```
int uv_fs_lchown(uv\_loop\_t\* loop, uv\_fs\_t\* req, const char* path, uv_uid_t uid, uv_gid_t gid, uv_fs_cb cb)
```

Equivalent to [chown\(2\)](#) [<https://man7.org/linux/man-pages/man2/chown.2.html>], [fchown\(2\)](#) [<https://man7.org/linux/man-pages/man2/fchown.2.html>] and [lchown\(2\)](#) [<https://man7.org/linux/man-pages/man2/lchown.2.html>] respectively.

## Note

These functions are not implemented on Windows.

*Changed in version 1.21.0:* implemented uv\_fs\_lchown

`uv_fs_type uv_fs_get_type(const uv\_fs\_t* req)`  
Returns `req->fs_type`.

*New in version 1.19.0.*

`ssize_t uv_fs_get_result(const uv\_fs\_t* req)`  
Returns `req->result`.

*New in version 1.19.0.*

`int uv_fs_get_system_error(const uv\_fs\_t* req)`  
Returns the platform specific error code - `GetLastError()` value on Windows and `-(req->result)` on other platforms.

*New in version 1.38.0.*

`void* uv_fs_get_ptr(const uv\_fs\_t* req)`  
Returns `req->ptr`.

*New in version 1.19.0.*

`const char* uv_fs_get_path(const uv\_fs\_t* req)`  
Returns `req->path`.

*New in version 1.19.0.*

`uv\_stat\_t* uv_fs_get_statbuf(uv\_fs\_t* req)`  
Returns `&req->statbuf`.

*New in version 1.19.0.*

## See also

The [uv\\_req\\_t](#) API functions also apply.



# Helper functions

[uv\\_os\\_fd\\_t](#) `uv_get_oshandle(int fd)`

For a file descriptor in the C runtime, get the OS-dependent handle. On UNIX, returns the `fd` intact. On Windows, this calls [\\_get\\_oshandle](#)

[<https://docs.microsoft.com/en-us/cpp/c-runtime-library/reference/get-oshandle?view=vs-2019>]. Note that the return value is still owned by the C runtime, any attempts to close it or to use it after closing the `fd` may lead to malfunction.

*New in version 1.12.0.*

`int uv_open_oshandle(uv\_os\_fd\_t os_fd)`

For a OS-dependent handle, get the file descriptor in the C runtime. On UNIX, returns the `os_fd` intact. On Windows, this calls [\\_open\\_oshandle](#) [<https://docs.microsoft.com/en-us/cpp/c-runtime-library/reference/open-oshandle?view=vs-2019>]. Note that this consumes the argument, any attempts to close it or to use it after closing the return value may lead to malfunction.

*New in version 1.23.0.*

## File open constants

**UV\_FS\_O\_APPEND**

The file is opened in append mode. Before each write, the file offset is positioned at the end of the file.

**UV\_FS\_O\_CREAT**

The file is created if it does not already exist.

**UV\_FS\_O\_DIRECT**

File I/O is done directly to and from user-space buffers, which must be aligned. Buffer size and address should be a multiple of the physical sector size of the block device.

**Note**

*UV\_FS\_O\_DIRECT* is supported on Linux, and on Windows via [FILE\\_FLAG\\_NO\\_BUFFERING](#)

[<https://docs.microsoft.com/en-us/windows/win32/fileio/file-buffering>].

*UV\_FS\_O\_DIRECT* is not supported on macOS.

**UV\_FS\_O\_DIRECTORY**

If the path is not a directory, fail the open.

**Note**

*UV\_FS\_O\_DIRECTORY* is not supported on Windows.

**UV\_FS\_O\_DSYNC**

The file is opened for synchronous I/O. Write operations will complete once all data and a minimum of metadata are flushed to disk.

**Note**

*UV\_FS\_O\_DSYNC* is supported on Windows via

[FILE\\_FLAG\\_WRITE\\_THROUGH](#) [<https://docs.microsoft.com/en-us/windows/win32/fileio/file-buffering>].

**UV\_FS\_O\_EXCL**

If the *O\_CREAT* flag is set and the file already exists, fail the open.

**Note**

In general, the behavior of `O_EXCL` is undefined if it is used without `O_CREAT`. There is one exception: on Linux 2.6 and later, `O_EXCL` can be used without `O_CREAT` if pathname refers to a block device. If the block device is in use by the system (e.g., mounted), the open will fail with the error `EBUSY`.

### **UV\_FS\_O\_EXLOCK**

Atomically obtain an exclusive lock.

#### **Note**

`UV_FS_O_EXLOCK` is only supported on macOS and Windows.

*Changed in version 1.17.0:* support is added for Windows.

### **UV\_FS\_O\_FILEMAP**

Use a memory file mapping to access the file. When using this flag, the file cannot be open multiple times concurrently.

#### **Note**

`UV_FS_O_FILEMAP` is only supported on Windows.

### **UV\_FS\_O\_NOATIME**

Do not update the file access time when the file is read.

#### **Note**

`UV_FS_O_NOATIME` is not supported on Windows.

### **UV\_FS\_O\_NOCTTY**

If the path identifies a terminal device, opening the path will not cause that terminal to become the controlling terminal for the process (if the process does not already have one).

**Note**

*UV\_FS\_O\_NOCTTY* is not supported on Windows.

**UV\_FS\_O\_NOFOLLOW**

If the path is a symbolic link, fail the open.

**Note**

*UV\_FS\_O\_NOFOLLOW* is not supported on Windows.

**UV\_FS\_O\_NONBLOCK**

Open the file in nonblocking mode if possible.

**Note**

*UV\_FS\_O\_NONBLOCK* is not supported on Windows.

**UV\_FS\_O\_RANDOM**

Access is intended to be random. The system can use this as a hint to optimize file caching.

**Note**

*UV\_FS\_O\_RANDOM* is only supported on Windows via [FILE\\_FLAG\\_RANDOM\\_ACCESS](https://docs.microsoft.com/en-us/windows/win32/api/fileapi/nf-fileapi-createfilea) [https://docs.microsoft.com/en-us/windows/win32/api/fileapi/nf-fileapi-createfilea].

**UV\_FS\_O\_RDONLY**

Open the file for read-only access.

#### **UV\_FS\_O\_RDWR**

Open the file for read-write access.

#### **UV\_FS\_O\_SEQUENTIAL**

Access is intended to be sequential from beginning to end. The system can use this as a hint to optimize file caching.

#### **Note**

*UV\_FS\_O\_SEQUENTIAL* is only supported on Windows via [FILE\\_FLAG\\_SEQUENTIAL\\_SCAN](https://docs.microsoft.com/en-us/windows/win32/api/fileapi/nf-fileapi-createfilea) [https://docs.microsoft.com/en-us/windows/win32/api/fileapi/nf-fileapi-createfilea].

#### **UV\_FS\_O\_SHORT\_LIVED**

The file is temporary and should not be flushed to disk if possible.

#### **Note**

*UV\_FS\_O\_SHORT\_LIVED* is only supported on Windows via [FILE\\_ATTRIBUTE\\_TEMPORARY](https://docs.microsoft.com/en-us/windows/win32/api/fileapi/nf-fileapi-createfilea) [https://docs.microsoft.com/en-us/windows/win32/api/fileapi/nf-fileapi-createfilea].

#### **UV\_FS\_O\_SYMLINK**

Open the symbolic link itself rather than the resource it points to.

#### **UV\_FS\_O\_SYNC**

The file is opened for synchronous I/O. Write operations will complete once all data and all metadata are flushed to disk.

**Note**

*UV\_FS\_O\_SYNC* is supported on Windows via [FILE\\_FLAG\\_WRITE\\_THROUGH](https://docs.microsoft.com/en-us/windows/win32/fileio/file-buffering) [https://docs.microsoft.com/en-us/windows/win32/fileio/file-buffering].

**UV\_FS\_O\_TEMPORARY**

The file is temporary and should not be flushed to disk if possible.

**Note**

*UV\_FS\_O\_TEMPORARY* is only supported on Windows via [FILE\\_ATTRIBUTE\\_TEMPORARY](https://docs.microsoft.com/en-us/windows/win32/api/fileapi/nf-fileapi-createfilea) [https://docs.microsoft.com/en-us/windows/win32/api/fileapi/nf-fileapi-createfilea].

**UV\_FS\_O\_TRUNC**

If the file exists and is a regular file, and the file is opened successfully for write access, its length shall be truncated to zero.

**UV\_FS\_O\_WRONLY**

Open the file for write-only access.

# Thread pool work scheduling

libuv provides a threadpool which can be used to run user code and get notified in the loop thread. This thread pool is internally used to run all file system operations, as well as getaddrinfo and getnameinfo requests.

Its default size is 4, but it can be changed at startup time by setting the `UV_THREADPOOL_SIZE` environment variable to any value (the absolute maximum is 1024).

*Changed in version 1.30.0:* the maximum `UV_THREADPOOL_SIZE` allowed was increased from 128 to 1024.

The threadpool is global and shared across all event loops. When a particular function makes use of the threadpool (i.e. when using [uv\\_queue\\_work\(\)](#)) libuv preallocates and initializes the maximum number of threads allowed by `UV_THREADPOOL_SIZE`. This causes a relatively minor memory overhead (~1MB for 128 threads) but increases the performance of threading at runtime.

## Note

Note that even though a global thread pool which is shared across all events loops is used, the functions are not thread safe.

## Data types

**uv\_work\_t**

Work request type.

void (\*uv\_work\_cb)([uv\\_work\\_t](#)\* req)

Callback passed to [uv\\_queue\\_work\(\)](#) which will be run on the thread pool.

void (\*uv\_after\_work\_cb)([uv\\_work\\_t](#)\* req, int status)

Callback passed to [uv\\_queue\\_work\(\)](#) which will be called on the loop thread after the work on the threadpool has been completed. If the work was cancelled using [uv\\_cancel\(\)](#) status will be UV\_ECANCELED.

## Public members

[uv\\_loop\\_t](#)\* uv\_work\_t.loop

Loop that started this request and where completion will be reported. Readonly.

## See also

The [uv\\_req\\_t](#) members also apply.

## API

int uv\_queue\_work([uv\\_loop\\_t](#)\* loop, [uv\\_work\\_t](#)\* req, [uv\\_work\\_cb](#) work\_cb, [uv\\_after\\_work\\_cb](#) after\_work\_cb)

Initializes a work request which will run the given *work\_cb* in a thread from the threadpool. Once *work\_cb* is completed, *after\_work\_cb* will be called on the loop thread.

This request can be cancelled with [uv\\_cancel\(\)](#).



## See also

The [uv\\_req\\_t](#) API functions also apply.

# DNS utility functions

libuv provides asynchronous variants of *getaddrinfo* and *getnameinfo*.

## Data types

**uv\_getaddrinfo\_t**

*getaddrinfo* request type.

void (\*[uv\\_getaddrinfo\\_cb](#))([uv\\_getaddrinfo\\_t](#)\* req, int status, struct addrinfo\* res)

Callback which will be called with the *getaddrinfo* request result once complete. In case it was cancelled, *status* will have a value of UV\_ECANCELED.

**uv\_getnameinfo\_t**

*getnameinfo* request type.

void (\*[uv\\_getnameinfo\\_cb](#))([uv\\_getnameinfo\\_t](#)\* req, int status, const char\* hostname, const char\* service)

Callback which will be called with the *getnameinfo* request result once complete. In case it was cancelled, *status* will have a value of UV\_ECANCELED.

## Public members

[uv\\_loop\\_t](#)\* **uv\_getaddrinfo\_t.loop**

Loop that started this *getaddrinfo* request and where completion will be reported. Readonly.

struct addrinfo\* **uv\_getaddrinfo\_t.addrinfo**

Pointer to a *struct addrinfo* containing the result. Must be freed by the user with [uv\\_freeaddrinfo\(\)](#).

*Changed in version 1.3.0:* the field is declared as public.

[uv\\_loop\\_t](#)\* **uv\_getnameinfo\_t.loop**

Loop that started this getnameinfo request and where completion will be reported. Readonly.

char[NI\_MAXHOST] **uv\_getnameinfo\_t.host**

Char array containing the resulting host. It's null terminated.

*Changed in version 1.3.0:* the field is declared as public.

char[NI\_MAXSERV] **uv\_getnameinfo\_t.service**

Char array containing the resulting service. It's null terminated.

*Changed in version 1.3.0:* the field is declared as public.

## See also

The [uv\\_req\\_t](#) members also apply.

## API

int **uv\_getaddrinfo**([uv\\_loop\\_t](#)\* loop, [uv\\_getaddrinfo\\_t](#)\* req, [uv\\_getaddrinfo\\_cb](#) getaddrinfo\_cb, const char\* node, const char\* service, const struct addrinfo\* hints)

Asynchronous [getaddrinfo\(3\)](#) [<https://man7.org/linux/man-pages/man3/getaddrinfo.3.html>].

Either node or service may be NULL but not both.

*hints* is a pointer to a struct `addrinfo` with additional address type constraints, or `NULL`. Consult *man -s 3 getaddrinfo* for more details.

Returns 0 on success or an error code < 0 on failure. If successful, the callback will get called sometime in the future with the lookup result, which is either:

- `status == 0`, the `res` argument points to a valid *struct addrinfo*, or
- `status < 0`, the `res` argument is `NULL`. See the `UV_EAI_*` constants.

Call [uv\\_freeaddrinfo\(\)](#) to free the `addrinfo` structure.

*Changed in version 1.3.0:* the callback parameter is now allowed to be `NULL`, in which case the request will run **synchronously**.

```
void uv_freeaddrinfo(struct addrinfo* ai)
```

Free the struct `addrinfo`. Passing `NULL` is allowed and is a no-op.

```
int uv_getnameinfo(uv_loop_t* loop, uv_getnameinfo_t* req,  
uv_getnameinfo_cb getnameinfo_cb, const struct  
sockaddr* addr, int flags)
```

Asynchronous [getnameinfo\(3\)](#) [<https://man7.org/linux/man-pages/man3/getnameinfo.3.html>].

Returns 0 on success or an error code < 0 on failure. If successful, the callback will get called sometime in the future with the lookup result. Consult *man -s 3 getnameinfo* for more details.

*Changed in version 1.3.0:* the callback parameter is now allowed to be `NULL`, in which case the request will run **synchronously**.

## **See also**

The [uv\\_req\\_t](#) API functions also apply.

# Shared library handling

libuv provides cross platform utilities for loading shared libraries and retrieving symbols from them, using the following API.

## Data types

`uv_lib_t`

Shared library data type.

## Public members

N/A

## API

`int uv_dlopen(const char* filename, uv\_lib\_t\* lib)`

Opens a shared library. The filename is in utf-8. Returns 0 on success and -1 on error. Call [uv\\_dlerror\(\)](#) to get the error message.

`void uv_dlclose(uv\_lib\_t\* lib)`

Close the shared library.

`int uv_dlsym(uv\_lib\_t\* lib, const char* name, void** ptr)`

Retrieves a data pointer from a dynamic library. It is legal for a symbol to map to NULL. Returns 0 on success and -1 if the symbol was not found.

`const char* uv_dlerror(const uv\_lib\_t\* lib)`

Returns the last `uv_dlopen()` or `uv_dlsym()` error message.

# Threading and synchronization utilities

libuv provides cross-platform implementations for multiple threading and synchronization primitives. The API largely follows the pthreads API.

## Data types

**uv\_thread\_t**

Thread data type.

**void (\*uv\_thread\_cb)(void\* arg)**

Callback that is invoked to initialize thread execution. *arg* is the same value that was passed to [uv\\_thread\\_create\(\)](#).

**uv\_key\_t**

Thread-local key data type.

**uv\_once\_t**

Once-only initializer data type.

**uv\_mutex\_t**

Mutex data type.

**uv\_rwlock\_t**

Read-write lock data type.

**uv\_sem\_t**

Semaphore data type.

**uv\_cond\_t**



Condition data type.

**uv\_barrier\_t**

Barrier data type.

## API

### Threads

**uv\_thread\_options\_t**

Options for spawning a new thread (passed to [uv\\_thread\\_create\\_ex\(\)](#)).

```
typedef struct uv_thread_options_s {
    enum {
        UV_THREAD_NO_FLAGS = 0x00,
        UV_THREAD_HAS_STACK_SIZE = 0x01
    } flags;
    size_t stack_size;
} uv_thread_options_t;
```

More fields may be added to this struct at any time, so its exact layout and size should not be relied upon.

*New in version 1.26.0.*

int **uv\_thread\_create**([uv\\_thread\\_t\\*](#) tid, [uv\\_thread\\_cb](#) entry, void\* arg)

*Changed in version 1.4.1:* returns a UV\_E\* error code on failure

int **uv\_thread\_create\_ex**([uv\\_thread\\_t\\*](#) tid, const [uv\\_thread\\_options\\_t\\*](#) params, [uv\\_thread\\_cb](#) entry, void\* arg)

Like [uv\\_thread\\_create\(\)](#), but additionally specifies options for creating a new thread.

If `UV_THREAD_HAS_STACK_SIZE` is set, `stack_size` specifies a stack size for the new thread. `0` indicates that the default value should be used, i.e. behaves as if the flag was not set. Other values will be rounded up to the nearest page boundary.

*New in version 1.26.0.*

```
uv\_thread\_t uv_thread_self(void)
```

```
int uv_thread_join(uv\_thread\_t *tid)
```

```
int uv_thread_equal(const uv\_thread\_t* t1, const  
uv\_thread\_t* t2)
```

## Thread-local storage

### Note

The total thread-local storage size may be limited. That is, it may not be possible to create many TLS keys.

```
int uv_key_create(uv\_key\_t* key)
```

```
void uv_key_delete(uv\_key\_t* key)
```

```
void* uv_key_get(uv\_key\_t* key)
```

```
void uv_key_set(uv\_key\_t* key, void* value)
```

## Once-only initialization

Runs a function once and only once. Concurrent calls to [uv\\_once\(\)](#) with the same guard will block all callers except

one (it's unspecified which one). The guard should be initialized statically with the `UV_ONCE_INIT` macro.

```
void uv_once(uv\_once\_t* guard, void (*callback)(void))
```

## Mutex locks

Functions return 0 on success or an error code < 0 (unless the return type is void, of course).

```
int uv_mutex_init(uv\_mutex\_t* handle)
```

```
int uv_mutex_init_recursive(uv\_mutex\_t* handle)
```

```
void uv_mutex_destroy(uv\_mutex\_t* handle)
```

```
void uv_mutex_lock(uv\_mutex\_t* handle)
```

```
int uv_mutex_trylock(uv\_mutex\_t* handle)
```

```
void uv_mutex_unlock(uv\_mutex\_t* handle)
```

## Read-write locks

Functions return 0 on success or an error code < 0 (unless the return type is void, of course).

```
int uv_rwlock_init(uv\_rwlock\_t* rwlock)
```

```
void uv_rwlock_destroy(uv\_rwlock\_t* rwlock)
```

```
void uv_rwlock_rdlock(uv\_rwlock\_t* rwlock)
```

```
int uv_rwlock_tryrdlock(uv\_rwlock\_t* rwlock)
```

```
void uv_rwlock_rdunlock(uv\_rwlock\_t* rwlock)
```

```
void uv_rwlock_wrlock(uv\_rwlock\_t* rwlock)
```

```
int uv_rwlock_trywrlock(uv\_rwlock\_t* rwlock)
```

```
void uv_rwlock_wrunlock(uv\_rwlock\_t* rwlock)
```

## Semaphores

Functions return 0 on success or an error code < 0 (unless the return type is void, of course).

```
int uv_sem_init(uv\_sem\_t* sem, unsigned int value)
```

```
void uv_sem_destroy(uv\_sem\_t* sem)
```

```
void uv_sem_post(uv\_sem\_t* sem)
```

```
void uv_sem_wait(uv\_sem\_t* sem)
```

```
int uv_sem_trywait(uv\_sem\_t* sem)
```

## Conditions

Functions return 0 on success or an error code < 0 (unless the return type is void, of course).

### Note

1. Callers should be prepared to deal with spurious wakeups on [uv\\_cond\\_wait\(\)](#) and [uv\\_cond\\_timedwait\(\)](#).
2. The timeout parameter for [uv\\_cond\\_timedwait\(\)](#) is relative to the time at which function is called.
3. On z/OS, the timeout parameter for [uv\\_cond\\_timedwait\(\)](#) is converted to an absolute system time at which the wait expires. If the current system clock time passes the absolute time calculated before the condition is

signaled, an ETIMEDOUT error results. After the wait begins, the wait time is not affected by changes to the system clock.

```
int uv_cond_init(uv\_cond\_t* cond)
```

```
void uv_cond_destroy(uv\_cond\_t* cond)
```

```
void uv_cond_signal(uv\_cond\_t* cond)
```

```
void uv_cond_broadcast(uv\_cond\_t* cond)
```

```
void uv_cond_wait(uv\_cond\_t* cond, uv\_mutex\_t* mutex)
```

```
int uv_cond_timedwait(uv\_cond\_t* cond, uv\_mutex\_t* mutex,  
uint64_t timeout)
```

## Barriers

Functions return 0 on success or an error code < 0 (unless the return type is void, of course).

### Note

[uv\\_barrier\\_wait\(\)](#) returns a value > 0 to an arbitrarily chosen “serializer” thread to facilitate cleanup, i.e.

```
if (uv_barrier_wait(&barrier) > 0)  
    uv_barrier_destroy(&barrier);
```

```
int uv_barrier_init(uv\_barrier\_t* barrier, unsigned  
int count)
```

```
void uv_barrier_destroy(uv\_barrier\_t* barrier)
```

```
int uv_barrier_wait(uv\_barrier\_t* barrier)
```

# Miscellaneous utilities

This section contains miscellaneous functions that don't really belong in any other section.

## Data types

**uv\_buf\_t**

Buffer data type.

**char\* uv\_buf\_t.base**

Pointer to the base of the buffer.

**size\_t uv\_buf\_t.len**

Total bytes in the buffer.

### Note

On Windows this field is ULONG.

**void\* (\*uv\_malloc\_func)(size\_t size)**

Replacement function for [malloc\(3\)](https://man7.org/linux/man-pages/man3/malloc.3.html) [https://man7.org/linux/man-pages/man3/malloc.3.html]. See [uv\\_replace\\_allocator\(\)](#).

**void\* (\*uv\_realloc\_func)(void\* ptr, size\_t size)**

Replacement function for [realloc\(3\)](https://man7.org/linux/man-pages/man3/realloc.3.html) [https://man7.org/linux/man-pages/man3/realloc.3.html]. See [uv\\_replace\\_allocator\(\)](#).

**void\* (\*uv\_calloc\_func)(size\_t count, size\_t size)**

Replacement function for [calloc\(3\)](https://man7.org/linux/man-pages/man3/calloc.3.html) [https://man7.org/linux/man-pages/man3/calloc.3.html]. See [uv\\_replace\\_allocator\(\)](#).

`void (*uv_free_func)(void* ptr)`

Replacement function for [free\(3\)](https://man7.org/linux/man-pages/man3/free.3.html) [https://man7.org/linux/man-pages/man3/free.3.html]. See [uv\\_replace\\_allocator\(\)](#).

`void (*uv_random_cb)(uv_random_t* req, int status, void* buf, size_t buflen)`

Callback passed to [uv\\_random\(\)](#). *status* is non-zero in case of error. The *buf* pointer is the same pointer that was passed to [uv\\_random\(\)](#).

**uv\_file**

Cross platform representation of a file handle.

**uv\_os\_sock\_t**

Cross platform representation of a socket handle.

**uv\_os\_fd\_t**

Abstract representation of a file descriptor. On Unix systems this is a *typedef* of *int* and on Windows a *HANDLE*.

**uv\_pid\_t**

Cross platform representation of a *pid\_t*.

*New in version 1.16.0.*

**uv\_timeval\_t**

Data type for storing times.

```
typedef struct {
    long tv_sec;
    long tv_usec;
} uv_timeval_t;
```

**uv\_timeval64\_t**

Alternative data type for storing times.



```
typedef struct {
    int64_t tv_sec;
    int32_t tv_usec;
} uv_timeval64_t;
```

## uv\_rusage\_t

Data type for resource usage results.

```
typedef struct {
    uv_timeval_t ru_utime; /* user CPU time used */
    uv_timeval_t ru_stime; /* system CPU time used */
    uint64_t ru_maxrss; /* maximum resident set size */
    uint64_t ru_ixrss; /* integral shared memory size (X) */
    uint64_t ru_idrss; /* integral unshared data size (X) */
    uint64_t ru_isrss; /* integral unshared stack size (X)
*/
    uint64_t ru_minflt; /* page reclaims (soft page faults)
(X) */
    uint64_t ru_majflt; /* page faults (hard page faults) */
    uint64_t ru_nswap; /* swaps (X) */
    uint64_t ru_inblock; /* block input operations */
    uint64_t ru_oublock; /* block output operations */
    uint64_t ru_msgsnd; /* IPC messages sent (X) */
    uint64_t ru_msgrcv; /* IPC messages received (X) */
    uint64_t ru_nsignals; /* signals received (X) */
    uint64_t ru_nvcsw; /* voluntary context switches (X) */
    uint64_t ru_nivcsw; /* involuntary context switches (X)
*/
} uv_rusage_t;
```

Members marked with (X) are unsupported on Windows.

See [getrusage\(2\)](https://man7.org/linux/man-pages/man2/getrusage.2.html) [https://man7.org/linux/man-

pages/man2/getrusage.2.html] for supported fields on Unix

## uv\_cpu\_info\_t

Data type for CPU information.

```
typedef struct uv_cpu_info_s {
    char* model;
    int speed;
    struct uv_cpu_times_s {
        uint64_t user; /* milliseconds */
        uint64_t nice; /* milliseconds */
    }
```

```
        uint64_t sys; /* milliseconds */
        uint64_t idle; /* milliseconds */
        uint64_t irq; /* milliseconds */
    } cpu_times;
} uv_cpu_info_t;
```

## **uv\_interface\_address\_t**

Data type for interface addresses.

```
typedef struct uv_interface_address_s {
    char* name;
    char phys_addr[6];
    int is_internal;
    union {
        struct sockaddr_in address4;
        struct sockaddr_in6 address6;
    } address;
    union {
        struct sockaddr_in netmask4;
        struct sockaddr_in6 netmask6;
    } netmask;
} uv_interface_address_t;
```

## **uv\_passwd\_t**

Data type for password file information.

```
typedef struct uv_passwd_s {
    char* username;
    long uid;
    long gid;
    char* shell;
    char* homedir;
} uv_passwd_t;
```

## **uv\_utsname\_t**

Data type for operating system name and version information.

```
typedef struct uv_utsname_s {
    char sysname[256];
    char release[256];
    char version[256];
}
```

```
    char machine[256];  
} uv_utsname_t;
```

## **uv\_env\_item\_t**

Data type for environment variable storage.

```
typedef struct uv_env_item_s {  
    char* name;  
    char* value;  
} uv_env_item_t;
```

## **uv\_random\_t**

Random data request type.

# **API**

## **uv\_handle\_type uv\_guess\_handle([uv\\_file](#) file)**

Used to detect what type of stream should be used with a given file descriptor. Usually this will be used during initialization to guess the type of the stdio streams.

For [isatty\(3\)](#) [<https://man7.org/linux/man-pages/man3/isatty.3.html>] equivalent functionality use this function and test for UV\_TTY.

## **int uv\_replace\_allocator([uv\\_malloc\\_func](#) malloc\_func, [uv\\_realloc\\_func](#) realloc\_func, [uv\\_calloc\\_func](#) calloc\_func, [uv\\_free\\_func](#) free\_func)**

*New in version 1.6.0.*

Override the use of the standard library's [malloc\(3\)](#) [<https://man7.org/linux/man-pages/man3/malloc.3.html>], [calloc\(3\)](#) [<https://man7.org/linux/man-pages/man3/calloc.3.html>], [realloc\(3\)](#) [<https://man7.org/linux/man-pages/man3/realloc.3.html>], [free\(3\)](#) [<https://man7.org/linux/man-pages/man3/free.3.html>], [memory](#) allocation functions.

This function must be called before any other libuv function is called or after all resources have been freed and thus libuv doesn't reference any allocated memory chunk.

On success, it returns 0, if any of the function pointers is NULL it returns UV\_EINVAL.

### **Warning**

There is no protection against changing the allocator multiple times. If the user changes it they are responsible for making sure the allocator is changed while no memory was allocated with the previous allocator, or that they are compatible.

### **Warning**

Allocator must be thread-safe.

**void uv\_library\_shutdown(void);**

*New in version 1.38.0.*

Release any global state that libuv is holding onto. Libuv will normally do so automatically when it is unloaded but it can be instructed to perform cleanup manually.

### **Warning**

Only call `uv_library_shutdown()` once.

### **Warning**

Don't call `uv_library_shutdown()` when there are still event loops or I/O requests active.

## Warning

Don't call libuv functions after calling `uv_library_shutdown()`.

[uv\\_buf\\_t](#) `uv_buf_init(char* base, unsigned int len)`

Constructor for [uv\\_buf\\_t](#).

Due to platform differences the user cannot rely on the ordering of the *base* and *len* members of the `uv_buf_t` struct. The user is responsible for freeing *base* after the `uv_buf_t` is done. Return struct passed by value.

`char** uv_setup_args(int argc, char** argv)`

Store the program arguments. Required for getting / setting the process title or the executable path. Libuv may take ownership of the memory that *argv* points to. This function should be called exactly once, at program start-up.

Example:

```
argv = uv_setup_args(argc, argv); /* May return a copy of  
argv. */
```

`int uv_get_process_title(char* buffer, size_t size)`

Gets the title of the current process. You *must* call `uv_setup_args` before calling this function on Unix and AIX systems. If `uv_setup_args` has not been called on systems that require it, then `UV_ENOBUFS` is returned. If *buffer* is `NULL` or *size* is zero, `UV_EINVAL` is returned. If *size* cannot accommodate the process title and terminating *nul* character, the function returns `UV_ENOBUFS`.

## Note

On BSD systems, *uv\_setup\_args* is needed for getting the initial process title. The process title returned will be an empty string until either *uv\_setup\_args* or *uv\_set\_process\_title* is called.

*Changed in version 1.18.1:* now thread-safe on all supported platforms.

*Changed in version 1.39.0:* now returns an error if *uv\_setup\_args* is needed but hasn't been called.

`int uv_set_process_title(const char* title)`

Sets the current process title. You *must* call *uv\_setup\_args* before calling this function on Unix and AIX systems. If *uv\_setup\_args* has not been called on systems that require it, then *UV\_ENOBUFS* is returned. On platforms with a fixed size buffer for the process title the contents of *title* will be copied to the buffer and truncated if larger than the available space. Other platforms will return *UV\_ENOMEM* if they cannot allocate enough space to duplicate the contents of *title*.

*Changed in version 1.18.1:* now thread-safe on all supported platforms.

*Changed in version 1.39.0:* now returns an error if *uv\_setup\_args* is needed but hasn't been called.

`int uv_resident_set_memory(size_t* rss)`

Gets the resident set size (RSS) for the current process.

`int uv_uptime(double* uptime)`

Gets the current system uptime. Depending on the system full or fractional seconds are returned.

int `uv_getrusage`([uv\\_rusage\\_t](#)\* *rusage*)

Gets the resource usage measures for the current process.

### **Note**

On Windows not all fields are set, the unsupported fields are filled with zeroes. See [uv\\_rusage\\_t](#) for more details.

[uv\\_pid\\_t](#) `uv_os_getpid`(void)

Returns the current process ID.

*New in version 1.18.0.*

[uv\\_pid\\_t](#) `uv_os_getppid`(void)

Returns the parent process ID.

*New in version 1.16.0.*

unsigned int `uv_available_parallelism`(void)

Returns an estimate of the default amount of parallelism a program should use. Always returns a non-zero value.

On Linux, inspects the calling thread's CPU affinity mask to determine if it has been pinned to specific CPUs.

On Windows, the available parallelism may be underreported on systems with more than 64 logical CPUs.

On other platforms, reports the number of CPUs that the operating system considers to be online.

*New in version 1.44.0.*

`int uv_cpu_info(uv\_cpu\_info\_t** cpu_infos, int* count)`  
Gets information about the CPUs on the system. The `cpu_infos` array will have `count` elements and needs to be freed with [uv\\_free\\_cpu\\_info\(\)](#).

Use [uv\\_available\\_parallelism\(\)](#) if you need to know how many CPUs are available for threads or child processes.

`void uv_free_cpu_info(uv\_cpu\_info\_t* cpu_infos, int count)`  
Frees the `cpu_infos` array previously allocated with [uv\\_cpu\\_info\(\)](#).

`int uv_interface_addresses(uv\_interface\_address\_t** addresses, int* count)`

Gets address information about the network interfaces on the system. An array of `count` elements is allocated and returned in `addresses`. It must be freed by the user, calling [uv\\_free\\_interface\\_addresses\(\)](#).

`void uv_free_interface_addresses(uv\_interface\_address\_t* addresses, int count)`

Free an array of [uv\\_interface\\_address\\_t](#) which was returned by [uv\\_interface\\_addresses\(\)](#).

`void uv_loadavg(double avg[3])`  
Gets the load average. See: [https://en.wikipedia.org/wiki/Load\\_\(computing\)](https://en.wikipedia.org/wiki/Load_(computing)).

### **Note**

Returns [0,0,0] on Windows (i.e., it's not implemented).



int **uv\_ip4\_addr**(const char\* *ip*, int *port*, struct sockaddr\_in\* *addr*)

Convert a string containing an IPv4 addresses to a binary structure.

int **uv\_ip6\_addr**(const char\* *ip*, int *port*, struct sockaddr\_in6\* *addr*)

Convert a string containing an IPv6 addresses to a binary structure.

int **uv\_ip4\_name**(const struct sockaddr\_in\* *src*, char\* *dst*, size\_t *size*)

Convert a binary structure containing an IPv4 address to a string.

int **uv\_ip6\_name**(const struct sockaddr\_in6\* *src*, char\* *dst*, size\_t *size*)

Convert a binary structure containing an IPv6 address to a string.

int **uv\_ip\_name**(const struct sockaddr \**src*, char \**dst*, size\_t *size*)

Convert a binary structure containing an IPv4 address or an IPv6 address to a string.

int **uv\_inet\_ntop**(int *af*, const void\* *src*, char\* *dst*, size\_t *size*)

int **uv\_inet\_pton**(int *af*, const char\* *src*, void\* *dst*)

Cross-platform IPv6-capable implementation of [inet\\_ntop\(3\)](https://man7.org/linux/man-pages/man3/inet_ntop.3.html) [https://man7.org/linux/man-pages/man3/inet\_ntop.3.html] and [inet\\_pton\(3\)](https://man7.org/linux/man-pages/man3/inet_pton.3.html) [https://man7.org/linux/man-pages/man3/inet\_pton.3.html]. On success they return 0. In case of error the target *dst* pointer is unmodified.

## UV\_IF\_NAMESIZE

Maximum IPv6 interface identifier name length. Defined as *IFNAMSIZ* on Unix and *IF\_NAMESIZE* on Linux and Windows.

*New in version 1.16.0.*

`int uv_if_indextoname(unsigned int ifindex, char* buffer, size_t* size)`

IPv6-capable implementation of [if\\_indextoname\(3\)](https://man7.org/linux/man-pages/man3/if_indextoname.3.html) [https://man7.org/linux/man-pages/man3/if\_indextoname.3.html]. When called, *\*size* indicates the length of the *buffer*, which is used to store the result. On success, zero is returned, *buffer* contains the interface name, and *\*size* represents the string length of the *buffer*, excluding the NUL terminator byte from *\*size*. On error, a negative result is returned. If *buffer* is not large enough to hold the result, *UV\_ENOBUFS* is returned, and *\*size* represents the necessary size in bytes, including the NUL terminator byte into the *\*size*.

On Unix, the returned interface name can be used directly as an interface identifier in scoped IPv6 addresses, e.g. *fe80::abc:def1:2345%en0*.

On Windows, the returned interface cannot be used as an interface identifier, as Windows uses numerical interface identifiers, e.g. *fe80::abc:def1:2345%5*.

To get an interface identifier in a cross-platform compatible way, use *uv\_if\_indextoiid()*.

Example:

```
char ifname[UV_IF_NAMESIZE];
size_t size = sizeof(ifname);
uv_if_indextoname(sin6->sin6_scope_id, ifname, &size);
```

*New in version 1.16.0.*

int **uv\_if\_index\_to\_id**(unsigned int *ifindex*, char\* *buffer*, size\_t\* *size*)

Retrieves a network interface identifier suitable for use in an IPv6 scoped address. On Windows, returns the numeric *ifindex* as a string. On all other platforms, *uv\_if\_index\_to\_name()* is called. The result is written to *buffer*, with *\*size* indicating the length of *buffer*. If *buffer* is not large enough to hold the result, then *UV\_ENOBUFS* is returned, and *\*size* represents the size, including the NUL byte, required to hold the result.

See *uv\_if\_index\_to\_name* for further details.

*New in version 1.16.0.*

int **uv\_exepath**(char\* *buffer*, size\_t\* *size*)

Gets the executable path. You *must* call *uv\_setup\_args* before calling this function.

int **uv\_cwd**(char\* *buffer*, size\_t\* *size*)

Gets the current working directory, and stores it in *buffer*. If the current working directory is too large to fit in *buffer*, this function returns *UV\_ENOBUFS*, and sets *size* to the required length, including the null terminator.

*Changed in version 1.1.0:* On Unix the path no longer ends in a slash.

*Changed in version 1.9.0:* the returned length includes the terminating null byte on *UV\_ENOBUFS*, and the buffer is null terminated on success.

int **uv\_chdir**(const char\* *dir*)

Changes the current working directory.

int `uv_os_homedir`(char\* *buffer*, size\_t\* *size*)

Gets the current user's home directory. On Windows, `uv_os_homedir()` first checks the `USERPROFILE` environment variable using `GetEnvironmentVariableW()`. If `USERPROFILE` is not set, `GetUserProfileDirectoryW()` is called. On all other operating systems, `uv_os_homedir()` first checks the `HOME` environment variable using [getenv\(3\)](https://man7.org/linux/man-pages/man3/getenv.3.html). [https://man7.org/linux/man-pages/man3/getenv.3.html]. If `HOME` is not set, [getpwuid\\_r\(3\)](https://man7.org/linux/man-pages/man3/getpwuid_r.3.html). [https://man7.org/linux/man-pages/man3/getpwuid\_r.3.html] is called. The user's home directory is stored in *buffer*. When `uv_os_homedir()` is called, *size* indicates the maximum size of *buffer*. On success *size* is set to the string length of *buffer*. On `UV_ENOBUFS` failure *size* is set to the required length for *buffer*, including the null byte.

### **Warning**

`uv_os_homedir()` is not thread safe.

*New in version 1.6.0.*

int `uv_os_tmpdir`(char\* *buffer*, size\_t\* *size*)

Gets the temp directory. On Windows, `uv_os_tmpdir()` uses `GetTempPathW()`. On all other operating systems, `uv_os_tmpdir()` uses the first environment variable found in the ordered list `TMPDIR`, `TMP`, `TEMP`, and `TEMPDIR`. If none of these are found, the path `"/tmp"` is used, or, on Android, `"/data/local/tmp"` is used. The temp directory is stored in *buffer*. When `uv_os_tmpdir()` is called, *size* indicates the maximum size of *buffer*. On success *size* is set to the string length of *buffer* (which does not include the terminating null). On `UV_ENOBUFS` failure *size* is set to the required length for *buffer*, including the null byte.

## Warning

`uv_os_tmpdir()` is not thread safe.

*New in version 1.9.0.*

`int uv_os_get_passwd(uv\_passwd\_t\* pwd)`

Gets a subset of the password file entry for the current effective uid (not the real uid). The populated data includes the username, euid, gid, shell, and home directory. On non-Windows systems, all data comes from [getpwuid\\_r\(3\)](#) [[https://man7.org/linux/man-pages/man3/getpwuid\\_r.3.html](https://man7.org/linux/man-pages/man3/getpwuid_r.3.html)]. On Windows, uid and gid are set to -1 and have no meaning, and shell is `NULL`. After successfully calling this function, the memory allocated to `pwd` needs to be freed with [uv\\_os\\_free\\_passwd\(\)](#).

*New in version 1.9.0.*

`void uv_os_free_passwd(uv\_passwd\_t\* pwd)`

Frees the `pwd` memory previously allocated with [uv\\_os\\_get\\_passwd\(\)](#).

*New in version 1.9.0.*

`uint64_t uv_get_free_memory(void)`

Gets the amount of free memory available in the system, as reported by the kernel (in bytes).

`uint64_t uv_get_total_memory(void)`

Gets the total amount of physical memory in the system (in bytes).

`uint64_t uv_get_constrained_memory(void)`

Gets the amount of memory available to the process (in bytes) based on limits imposed by the OS. If there is no such constraint, or the constraint is unknown, 0 is returned. Note that it is not unusual for this value to be less than or greater than [uv\\_get\\_total\\_memory\(\)](#).

### **Note**

This function currently only returns a non-zero value on Linux, based on cgroups if it is present, and on z/OS based on RLIMIT\_MEMLIMIT.

*New in version 1.29.0.*

`uint64_t uv_hrtime(void)`

Returns the current high-resolution real time. This is expressed in nanoseconds. It is relative to an arbitrary time in the past. It is not related to the time of day and therefore not subject to clock drift. The primary use is for measuring performance between intervals.

### **Note**

Not every platform can support nanosecond resolution; however, this value will always be in nanoseconds.

`void uv_print_all_handles(uv_loop_t* loop, FILE* stream)`

Prints all handles associated with the given *loop* to the given *stream*.

Example:

```
uv_print_all_handles(uv_default_loop(), stderr);
/*
[ -I] signal    0x1a25ea8
[ -AI] async   0x1a25cf0
```

```
[R--] idle      0x1a7a8c8
*/
```

The format is *[flags] handle-type handle-address*. For *flags*:

- *R* is printed for a handle that is referenced
- *A* is printed for a handle that is active
- *I* is printed for a handle that is internal

### **Warning**

This function is meant for ad hoc debugging, there is no API/ABI stability guarantees.

*New in version 1.8.0.*

```
void uv_print_active_handles(uv\_loop\_t* loop, FILE* stream)
```

This is the same as [uv\\_print\\_all\\_handles\(\)](#) except only active handles are printed.

### **Warning**

This function is meant for ad hoc debugging, there is no API/ABI stability guarantees.

*New in version 1.8.0.*

```
int uv_os_environ(uv\_env\_item\_t** envitems, int* count)
```

Retrieves all environment variables. This function will allocate memory which must be freed by calling [uv\\_os\\_free\\_environ\(\)](#).

### **Warning**

This function is not thread safe.

*New in version 1.31.0.*

```
void uv_os_free_envIRON(uv_env_item_t* envitems, int count);
```

Frees the memory allocated for the environment variables by [uv\\_os\\_envIRON\(\)](#).

*New in version 1.31.0.*

```
int uv_os_getenv(const char* name, char* buffer, size_t* size)
```

Retrieves the environment variable specified by *name*, copies its value into *buffer*, and sets *size* to the string length of the value. When calling this function, *size* must be set to the amount of storage available in *buffer*, including the null terminator. If the environment variable exceeds the storage available in *buffer*, *UV\_ENOBUFS* is returned, and *size* is set to the amount of storage required to hold the value. If no matching environment variable exists, *UV\_ENOENT* is returned.

### **Warning**

This function is not thread safe.

*New in version 1.12.0.*

```
int uv_os_setenv(const char* name, const char* value)
```

Creates or updates the environment variable specified by *name* with *value*.

### **Warning**

This function is not thread safe.

*New in version 1.12.0.*



int `uv_os_unsetenv`(const char\* *name*)

Deletes the environment variable specified by *name*. If no such environment variable exists, this function returns successfully.

### **Warning**

This function is not thread safe.

*New in version 1.12.0.*

int `uv_os_gethostname`(char\* *buffer*, size\_t\* *size*)

Returns the hostname as a null-terminated string in *buffer*, and sets *size* to the string length of the hostname. When calling this function, *size* must be set to the amount of storage available in *buffer*, including the null terminator. If the hostname exceeds the storage available in *buffer*, `UV_ENOBUFS` is returned, and *size* is set to the amount of storage required to hold the value.

*New in version 1.12.0.*

*Changed in version 1.26.0: `UV_MAXHOSTNAMESIZE` is available and represents the maximum *buffer* size required to store a hostname and terminating *nul* character.*

int `uv_os_getpriority`([uv\\_pid\\_t](#) *pid*, int\* *priority*)

Retrieves the scheduling priority of the process specified by *pid*. The returned value of *priority* is between -20 (high priority) and 19 (low priority).

### **Note**

On Windows, the returned priority will equal one of the `UV_PRIORITY` constants.

*New in version 1.23.0.*

int `uv_os_setpriority`([uv\\_pid\\_t](#) *pid*, int *priority*)

Sets the scheduling priority of the process specified by *pid*. The *priority* value range is between -20 (high priority) and 19 (low priority). The constants `UV_PRIORITY_LOW`, `UV_PRIORITY_BELOW_NORMAL`, `UV_PRIORITY_NORMAL`, `UV_PRIORITY_ABOVE_NORMAL`, `UV_PRIORITY_HIGH`, and `UV_PRIORITY_HIGHEST` are also provided for convenience.

### **Note**

On Windows, this function utilizes `SetPriorityClass()`. The *priority* argument is mapped to a Windows priority class. When retrieving the process priority, the result will equal one of the `UV_PRIORITY` constants, and not necessarily the exact value of *priority*.

### **Note**

On Windows, setting `PRIORITY_HIGHEST` will only work for elevated user, for others it will be silently reduced to `PRIORITY_HIGH`.

### **Note**

On IBM i PASE, the highest process priority is -10. The constant `UV_PRIORITY_HIGHEST` is -10, `UV_PRIORITY_HIGH` is -7, `UV_PRIORITY_ABOVE_NORMAL` is -4, `UV_PRIORITY_NORMAL` is 0, `UV_PRIORITY_BELOW_NORMAL` is 15 and `UV_PRIORITY_LOW` is 39.

## Note

On IBM i PASE, you are not allowed to change your priority unless you have the \*JOBCTL special authority (even to lower it).

*New in version 1.23.0.*

int **uv\_os\_uname**([uv\\_utsname\\_t](#)\* *buffer*)

Retrieves system information in *buffer*. The populated data includes the operating system name, release, version, and machine. On non-Windows systems, *uv\_os\_uname()* is a thin wrapper around [uname\(2\)](#). [https://man7.org/linux/man-pages/man2/uname.2.html]. Returns zero on success, and a non-zero error value otherwise.

*New in version 1.25.0.*

int **uv\_gettimeofday**([uv\\_timeval64\\_t](#)\* *tv*)

Cross-platform implementation of [gettimeofday\(2\)](#). [https://man7.org/linux/man-pages/man2/gettimeofday.2.html]. The timezone argument to *gettimeofday()* is not supported, as it is considered obsolete.

*New in version 1.28.0.*

int **uv\_random**([uv\\_loop\\_t](#)\* *loop*, [uv\\_random\\_t](#)\* *req*, void\* *buf*, [size\\_t](#) *buflen*, unsigned int *flags*, [uv\\_random\\_cb](#) *cb*)

Fill *buf* with exactly *buflen* cryptographically strong random bytes acquired from the system CSPRNG. *flags* is reserved for future extension and must currently be 0.

Short reads are not possible. When less than *buflen* random bytes are available, a non-zero error value is returned or passed to the callback.

The synchronous version may block indefinitely when not enough entropy is available. The asynchronous version may not ever finish when the system is low on entropy.

Sources of entropy:

- Windows: *RtlGenRandom*  
<<https://docs.microsoft.com/en-us/windows/desktop/api/ntsecapi/nf-ntsecapi-rtlgenrandom>>\_.
- Linux, Android: [getrandom\(2\)](https://man7.org/linux/man-pages/man2/getrandom.2.html) [https://man7.org/linux/man-pages/man2/getrandom.2.html] if available, or [urandom\(4\)](https://man7.org/linux/man-pages/man4/urandom.4.html) [https://man7.org/linux/man-pages/man4/urandom.4.html] after reading from */dev/random* once, or the *KERN\_RANDOM* [sysctl\(2\)](https://man7.org/linux/man-pages/man2/sysctl.2.html) [https://man7.org/linux/man-pages/man2/sysctl.2.html].
- FreeBSD: *getrandom(2)*  
<<https://www.freebsd.org/cgi/man.cgi?query=getrandom&sektion=2>>\_, or */dev/urandom* after reading from */dev/random* once.
- NetBSD: *KERN\_ARND* *sysctl(7)*  
<<https://man.netbsd.org/sysctl.7>>\_.
- macOS, OpenBSD: *getentropy(2)*  
<<https://man.openbsd.org/getentropy.2>>\_ if available, or */dev/urandom* after reading from */dev/random* once.
- AIX: */dev/random*.
- IBM i: */dev/urandom*.
- Other UNIX: */dev/urandom* after reading from */dev/random* once.

**Returns:** 0 on success, or an error code < 0 on failure. The contents of *buf* is undefined after an error.

**Note**

When using the synchronous version, both *loop* and *req* parameters are not used and can be set to *NULL*.

*New in version 1.33.0.*

void **uv\_sleep**(unsigned int *msec*)

Causes the calling thread to sleep for *msec* milliseconds.

*New in version 1.34.0.*

# Metrics operations

libuv provides a metrics API to track the amount of time the event loop has spent idle in the kernel's event provider.

## API

`uint64_t uv_metrics_idle_time(uv\_loop\_t\* loop)`

Retrieve the amount of time the event loop has been idle in the kernel's event provider (e.g. `epoll_wait`). The call is thread safe.

The return value is the accumulated time spent idle in the kernel's event provider starting from when the [uv\\_loop\\_t](#) was configured to collect the idle time.

### Note

The event loop will not begin accumulating the event provider's idle time until calling [uv\\_loop\\_configure](#) with `UV_METRICS_IDLE_TIME`.

*New in version 1.39.0.*

# User guide

## Warning

The contents of this guide have been recently incorporated into the libuv documentation and it hasn't gone through thorough review yet. If you spot a mistake please file an issue, or better yet, open a pull request!

- [Introduction](#)
  - [Who this book is for](#)
  - [Background](#)
  - [Code](#)
- [Basics of libuv](#)
  - [Event loops](#)
  - [Hello World](#)
  - [Error handling](#)
  - [Handles and Requests](#)
- [Filesystem](#)
  - [Reading/Writing files](#)
  - [Filesystem operations](#)
  - [Buffers and Streams](#)
  - [File change events](#)
- [Networking](#)
  - [TCP](#)
  - [UDP](#)
  - [Querying DNS](#)
  - [Network interfaces](#)
- [Threads](#)
  - [Core thread operations](#)
  - [Synchronization Primitives](#)
  - [libuv work queue](#)

- [Inter-thread communication](#)
- [Processes](#)
  - [Spawning child processes](#)
  - [Changing process parameters](#)
  - [Detaching processes](#)
  - [Sending signals to processes](#)
  - [Signals](#)
  - [Child Process I/O](#)
  - [Parent-child IPC](#)
- [Advanced event loops](#)
  - [Stopping an event loop](#)
- [Utilities](#)
  - [Timers](#)
  - [Event loop reference count](#)
  - [Idler pattern](#)
  - [Passing data to worker thread](#)
  - [External I/O with polling](#)
  - [Check & Prepare watchers](#)
  - [Loading libraries](#)
  - [TTY](#)
- [About](#)



# Introduction

This 'book' is a small set of tutorials about using [libuv](https://github.com/libuv/libuv) [https://github.com/libuv/libuv] as a high performance evented I/O library which offers the same API on Windows and Unix.

It is meant to cover the main areas of libuv, but is not a comprehensive reference discussing every function and data structure. The [official libuv documentation](https://docs.libuv.org/en/v1.x/) [https://docs.libuv.org/en/v1.x/] may be consulted for full details.

This book is still a work in progress, so sections may be incomplete, but I hope you will enjoy it as it grows.

# Who this book is for

If you are reading this book, you are either:

1. a systems programmer, creating low-level programs such as daemons or network services and clients. You have found that the event loop approach is well suited for your application and decided to use libuv.
2. a node.js module writer, who wants to wrap platform APIs written in C or C++ with a set of (a)synchronous APIs that are exposed to JavaScript. You will use libuv purely in the context of node.js. For this you will require some other resources as the book does not cover parts specific to v8/node.js.

This book assumes that you are comfortable with the C programming language.

## Background

The [node.js](https://www.nodejs.org) [https://www.nodejs.org] project began in 2009 as a JavaScript environment decoupled from the browser. Using Google's [V8](https://v8.dev) [https://v8.dev] and Marc Lehmann's [libev](http://software.schmorp.de/pkg/libev.html) [http://software.schmorp.de/pkg/libev.html], node.js combined a model of I/O - evented - with a language that was well suited to the style of programming; due to the way it had been shaped by browsers. As node.js grew in popularity, it was important to make it work on Windows, but libev ran only on Unix. The Windows equivalent of kernel event notification mechanisms like kqueue or (e)poll is IOCP. libuv was an abstraction around libev or IOCP depending on the platform, providing users an API based on libev. In the node-v0.9.0 version of libuv [libev was removed](https://github.com/joyent/libuv/issues/485) [https://github.com/joyent/libuv/issues/485].

Since then libuv has continued to mature and become a high quality standalone library for system programming. Users outside of node.js include Mozilla's [Rust](https://www.rust-lang.org) programming language, and a [variety](https://github.com/libuv/libuv/blob/v1.x/LINKS.md) of language bindings.

This book and the code is based on libuv version [v1.42.0](https://github.com/libuv/libuv/releases/tag/v1.42.0) [https://github.com/libuv/libuv/releases/tag/v1.42.0].

## Code

All the example code and the source of the book is included as part of the [libuv](https://github.com/libuv/libuv) project on GitHub. Clone or Download [libuv](https://github.com/libuv/libuv), then build it:

```
sh autogen.sh
./configure
make
```

There is no need to make `install`. To build the examples run `make` in the `docs/code/` directory.

# Basics of libuv

libuv enforces an **asynchronous, event-driven** style of programming. Its core job is to provide an event loop and callback based notifications of I/O and other activities. libuv offers core utilities like timers, non-blocking networking support, asynchronous file system access, child processes and more.

## Event loops

In event-driven programming, an application expresses interest in certain events and respond to them when they occur. The responsibility of gathering events from the operating system or monitoring other sources of events is handled by libuv, and the user can register callbacks to be invoked when an event occurs. The event-loop usually keeps running *forever*. In pseudocode:

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

Some examples of events are:

- File is ready for writing
- A socket has data ready to be read
- A timer has timed out

This event loop is encapsulated by `uv_run()` – the end-all function when using libuv.

The most common activity of systems programs is to deal with input and output, rather than a lot of number-

crunching. The problem with using conventional input/output functions (`read`, `fprintf`, etc.) is that they are **blocking**. The actual write to a hard disk or reading from a network, takes a disproportionately long time compared to the speed of the processor. The functions don't return until the task is done, so that your program is doing nothing. For programs which require high performance this is a major roadblock as other activities and other I/O operations are kept waiting.

One of the standard solutions is to use threads. Each blocking I/O operation is started in a separate thread (or in a thread pool). When the blocking function gets invoked in the thread, the operating system can schedule another thread to run, which actually needs the CPU.


The approach followed by libuv uses another style, which is the **asynchronous, non-blocking** style. Most modern operating systems provide event notification subsystems. For example, a normal `read` call on a socket would block until the sender actually sent something. Instead, the application can request the operating system to watch the socket and put an event notification in the queue. The application can inspect the events at its convenience (perhaps doing some number crunching before to use the processor to the maximum) and grab the data. It is **asynchronous** because the application expressed interest at one point, then used the data at another point (in time and space). It is **non-blocking** because the application process was free to do other tasks. This fits in well with libuv's event-loop approach, since the operating system events can be treated as just another libuv event. The non-blocking ensures that other events can continue to be handled as fast as they come in [\[1\]](#).

## Note

How the I/O is run in the background is not of our concern, but due to the way our computer hardware works, with the thread as the basic unit of the processor, libuv and OSes will usually run background/worker threads and/or polling to perform tasks in a non-blocking manner.

Bert Belder, one of the libuv core developers has a small video explaining the architecture of libuv and its background. If you have no prior experience with either libuv or libev, it is a quick, useful watch.

libuv's event loop is explained in more detail in the [documentation](https://docs.libuv.org/en/v1.x/design.html#the-i-o-loop) [https://docs.libuv.org/en/v1.x/design.html#the-i-o-loop].



## Hello World

With the basics out of the way, let's write our first libuv program. It does nothing, except start a loop which will exit immediately.

## helloworld/main.c

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <uv.h>
4
5  int main() {
6      uv_loop_t *loop = malloc(sizeof(uv_loop_t));
7      uv_loop_init(loop);
8
9      printf("Now quitting.\n");
10     uv_run(loop, UV_RUN_DEFAULT);
11
12     uv_loop_close(loop);
13     free(loop);
14     return 0;
15 }
```

This program quits immediately because it has no events to process. A libuv event loop has to be told to watch out for events using the various API functions.

Starting with libuv v1.0, users should allocate the memory for the loops before initializing it with `uv_loop_init(uv_loop_t *)`. This allows you to plug in custom memory management. Remember to de-initialize the loop using `uv_loop_close(uv_loop_t *)` and then delete the storage. The examples never close loops since the program quits after the loop ends and the system will reclaim memory. Production grade projects, especially long running systems programs, should take care to release correctly.

## Default loop

A default loop is provided by libuv and can be accessed using `uv_default_loop()`. You should use this loop if you only want a single loop.

## default-loop/main.c

```
1  #include <stdio.h>
2  #include <uv.h>
3
4  int main() {
5      uv_loop_t *loop = uv_default_loop();
6
7      printf("Default loop.\n");
8      uv_run(loop, UV_RUN_DEFAULT);
9
10     uv_loop_close(loop);
11     return 0;
12 }
```

### Note

node.js uses the default loop as its main loop. If you are writing bindings you should be aware of this.

## Error handling

Initialization functions or synchronous functions which may fail return a negative number on error. Async functions that may fail will pass a status parameter to their callbacks. The error messages are defined as UV\_E\* [constants](#)

[<https://docs.libuv.org/en/v1.x/errors.html#error-constants>].

You can use the `uv_strerror(int)` and `uv_err_name(int)` functions to get a `const char *` describing the error or the error name respectively.

I/O read callbacks (such as for files and sockets) are passed a parameter `nread`. If `nread` is less than 0, there was an error (UV\_EOF is the end of file error, which you may want to handle differently).



# Handles and Requests

libuv works by the user expressing interest in particular events. This is usually done by creating a **handle** to an I/O device, timer or process. Handles are opaque structs named as `uv_TYPE_t` where type signifies what the handle is used for.

## libuv watchers

```
/* Handle types. */
typedef struct uv_loop_s uv_loop_t;
typedef struct uv_handle_s uv_handle_t;
typedef struct uv_dir_s uv_dir_t;
typedef struct uv_stream_s uv_stream_t;
typedef struct uv_tcp_s uv_tcp_t;
typedef struct uv_udp_s uv_udp_t;
typedef struct uv_pipe_s uv_pipe_t;
typedef struct uv_tty_s uv_tty_t;
typedef struct uv_poll_s uv_poll_t;
typedef struct uv_timer_s uv_timer_t;
typedef struct uv_prepare_s uv_prepare_t;
typedef struct uv_check_s uv_check_t;
typedef struct uv_idle_s uv_idle_t;
typedef struct uv_async_s uv_async_t;
typedef struct uv_process_s uv_process_t;
typedef struct uv_fs_event_s uv_fs_event_t;
typedef struct uv_fs_poll_s uv_fs_poll_t;
typedef struct uv_signal_s uv_signal_t;

/* Request types. */
typedef struct uv_req_s uv_req_t;
typedef struct uv_getaddrinfo_s uv_getaddrinfo_t;
typedef struct uv_getnameinfo_s uv_getnameinfo_t;
typedef struct uv_shutdown_s uv_shutdown_t;
typedef struct uv_write_s uv_write_t;
typedef struct uv_connect_s uv_connect_t;
typedef struct uv_udp_send_s uv_udp_send_t;
typedef struct uv_fs_s uv_fs_t;
typedef struct uv_work_s uv_work_t;
typedef struct uv_random_s uv_random_t;

/* None of the above. */
```

```
typedef struct uv_env_item_s uv_env_item_t;
typedef struct uv_cpu_info_s uv_cpu_info_t;
typedef struct uv_interface_address_s uv_interface_address_t;
typedef struct uv_dirent_s uv_dirent_t;
typedef struct uv_passwd_s uv_passwd_t;
typedef struct uv_utsname_s uv_utsname_t;
typedef struct uv_statfs_s uv_statfs_t;
```

Handles represent long-lived objects. Async operations on such handles are identified using **requests**. A request is short-lived (usually used across only one callback) and usually indicates one I/O operation on a handle. Requests are used to preserve context between the initiation and the callback of individual actions. For example, an UDP socket is represented by a `uv_udp_t`, while individual writes to the socket use a `uv_udp_send_t` structure that is passed to the callback after the write is done.

Handles are setup by a corresponding:

```
uv_TYPE_init(uv_loop_t *, uv_TYPE_t *)
```

function.

Callbacks are functions which are called by libuv whenever an event the watcher is interested in has taken place. Application specific logic will usually be implemented in the callback. For example, an IO watcher's callback will receive the data read from a file, a timer callback will be triggered on timeout and so on.

## Idling

Here is an example of using an idle handle. The callback is called once on every turn of the event loop. A use case for idle handles is discussed in [Utilities](#). Let us use an idle watcher to look at the watcher life cycle and see how `uv_run()` will now block because a watcher is present. The

idle watcher is stopped when the count is reached and `uv_run()` exits since no event watchers are active.

## idle-basic/main.c

```
#include <stdio.h>
#include <uv.h>

int64_t counter = 0;

void wait_for_a_while(uv_idle_t* handle) {
    counter++;

    if (counter >= 10e6)
        uv_idle_stop(handle);
}

int main() {
    uv_idle_t idler;

    uv_idle_init(uv_default_loop(), &idler);
    uv_idle_start(&idler, wait_for_a_while);

    printf("Idling...\n");
    uv_run(uv_default_loop(), UV_RUN_DEFAULT);

    uv_loop_close(uv_default_loop());
    return 0;
}
```

## Storing context

In callback based programming style you'll often want to pass some 'context' - application specific information - between the call site and the callback. All handles and requests have a `void*` data member which you can set to the context and cast back in the callback. This is a common pattern used throughout the C library ecosystem. In addition `uv_loop_t` also has a similar data member.

---

[[1](#)] Depending on the capacity of the hardware of course.

# Filesystem

Simple filesystem read/write is achieved using the `uv_fs_*` functions and the `uv_fs_t` struct.

## Note

The libuv filesystem operations are different from [socket operations](#). Socket operations use the non-blocking operations provided by the operating system. Filesystem operations use blocking functions internally, but invoke these functions in a [thread pool](#)

[<https://docs.libuv.org/en/v1.x/threadpool.html#thread-pool-work-scheduling>] and notify watchers registered with the event loop when application interaction is required.

All filesystem functions have two forms - *synchronous* and *asynchronous*.

The *synchronous* forms automatically get called (and **block**) if the callback is null. The return value of functions is a [libuv error code](#). This is usually only useful for synchronous calls. The *asynchronous* form is called when a callback is passed and the return value is 0.

## Reading/Writing files

A file descriptor is obtained using

```
int uv_fs_open(uv_loop_t* loop, uv_fs_t* req, const char* path,
int flags, int mode, uv_fs_cb cb)
```

flags and mode are standard [Unix flags](https://man7.org/linux/man-pages/man2/open.2.html) [https://man7.org/linux/man-pages/man2/open.2.html]. libuv takes care of converting to the appropriate Windows flags.

File descriptors are closed using

```
int uv_fs_close(uv_loop_t* loop, uv_fs_t* req, uv_file file, uv_fs_cb cb)
```

Filesystem operation callbacks have the signature:

```
void callback(uv_fs_t* req);
```

Let's see a simple implementation of cat. We start with registering a callback for when the file is opened:

### uvcat/main.c - opening a file

```
1 // The request passed to the callback is the same
2 // as the one the call setup
3 // function was passed.
4 assert(req == &open_req);
5 if (req->result >= 0) {
6     iov = uv_buf_init(buffer, sizeof(buffer));
7     uv_fs_read(uv_default_loop(), &read_req, req->
8 >result,
9                 &iov, 1, -1, on_read);
10 }
11 else {
12     fprintf(stderr, "error opening file: %s\n",
13 uv_strerror((int)req->result));
14 }
```

The result field of a uv\_fs\_t is the file descriptor in case of the uv\_fs\_open callback. If the file is successfully opened, we start reading it.

## uvcat/main.c - read callback

```
1     if (req->result < 0) {
2         fprintf(stderr, "Read error: %s\n",
3             uv_strerror(req->result));
4     }
5     else if (req->result == 0) {
6         uv_fs_t close_req;
7         // synchronous
8         uv_fs_close(uv_default_loop(), &close_req,
9             open_req.result, NULL);
10    }
11    else if (req->result > 0) {
12        iov.len = req->result;
13        uv_fs_write(uv_default_loop(), &write_req, 1,
14            &iov, 1, -1, on_write);
15    }
16 }
```

In the case of a read call, you should pass an *initialized* buffer which will be filled with data before the read callback is triggered. The `uv_fs_*` operations map almost directly to certain POSIX functions, so EOF is indicated in this case by `result` being 0. In the case of streams or pipes, the `UV_EOF` constant would have been passed as a status instead.

Here you see a common pattern when writing asynchronous programs. The `uv_fs_close()` call is performed synchronously. *Usually tasks which are one-off, or are done as part of the startup or shutdown stage are performed synchronously, since we are interested in fast I/O when the program is going about its primary task and dealing with multiple I/O sources.* For solo tasks the performance difference usually is negligible and may lead to simpler code.

Filesystem writing is similarly simple using `uv_fs_write()`. *Your callback will be triggered after the write is complete.* In

our case the callback simply drives the next read. Thus read and write proceed in lockstep via callbacks.

## uvcat/main.c - write callback

```
1     if (req->result < 0) {
2         fprintf(stderr, "Write error: %s\n",
3             uv_strerror((int)req->result));
4     }
5     else {
6         uv_fs_read(uv_default_loop(), &read_req,
7             open_req.result, &iiov, 1, -1, on_read);
8     }
9 }
```

## Warning

Due to the way filesystems and disk drives are configured for performance, a write that 'succeeds' may not be committed to disk yet.

We set the dominos rolling in `main()`:

## uvcat/main.c

```
1     uv_fs_open(uv_default_loop(), &open_req, argv[1],
2         O_RDONLY, 0, on_open);
3     uv_run(uv_default_loop(), UV_RUN_DEFAULT);
4
5     uv_fs_req_cleanup(&open_req);
6     uv_fs_req_cleanup(&read_req);
7     uv_fs_req_cleanup(&write_req);
8     return 0;
9 }
```



## Warning

The `uv_fs_req_cleanup()` function must always be called on filesystem requests to free internal memory allocations in libuv.

## Filesystem operations

All the standard filesystem operations like `unlink`, `rmdir`, `stat` are supported asynchronously and have intuitive argument order. They follow the same patterns as the `read/write/open` calls, returning the result in the `uv_fs_t.result` field. The full list:

### Filesystem operations

```
int uv_fs_close(uv_loop_t* loop, uv_fs_t* req, uv_file file,
uv_fs_cb cb);
int uv_fs_open(uv_loop_t* loop, uv_fs_t* req, const char* path,
int flags, int mode, uv_fs_cb cb);
int uv_fs_read(uv_loop_t* loop, uv_fs_t* req, uv_file file,
const uv_buf_t bufs[], unsigned int nbufs, int64_t offset,
uv_fs_cb cb);
int uv_fs_unlink(uv_loop_t* loop, uv_fs_t* req, const char*
path, uv_fs_cb cb);
int uv_fs_write(uv_loop_t* loop, uv_fs_t* req, uv_file file,
const uv_buf_t bufs[], unsigned int nbufs, int64_t offset,
uv_fs_cb cb);
int uv_fs_copyfile(uv_loop_t* loop, uv_fs_t* req, const char*
path, const char* new_path, int flags, uv_fs_cb cb);
int uv_fs_mkdir(uv_loop_t* loop, uv_fs_t* req, const char*
path, int mode, uv_fs_cb cb);
int uv_fs_mkdtemp(uv_loop_t* loop, uv_fs_t* req, const char*
tpl, uv_fs_cb cb);
int uv_fs_mkstemp(uv_loop_t* loop, uv_fs_t* req, const char*
tpl, uv_fs_cb cb);
int uv_fs_rmdir(uv_loop_t* loop, uv_fs_t* req, const char*
path, uv_fs_cb cb);
int uv_fs_scandir(uv_loop_t* loop, uv_fs_t* req, const char*
```

```
path, int flags, uv_fs_cb cb);
int uv_fs_scandir_next(uv_fs_t* req, uv_dirent_t* ent);
int uv_fs_opendir(uv_loop_t* loop, uv_fs_t* req, const char*
path, uv_fs_cb cb);
int uv_fs_readdir(uv_loop_t* loop, uv_fs_t* req, uv_dir_t* dir,
uv_fs_cb cb);
int uv_fs_closedir(uv_loop_t* loop, uv_fs_t* req, uv_dir_t*
dir, uv_fs_cb cb);
int uv_fs_stat(uv_loop_t* loop, uv_fs_t* req, const char* path,
uv_fs_cb cb);
int uv_fs_fstat(uv_loop_t* loop, uv_fs_t* req, uv_file file,
uv_fs_cb cb);
int uv_fs_rename(uv_loop_t* loop, uv_fs_t* req, const char*
path, const char* new_path, uv_fs_cb cb);
int uv_fs_fsync(uv_loop_t* loop, uv_fs_t* req, uv_file file,
uv_fs_cb cb);
int uv_fs_fdatasync(uv_loop_t* loop, uv_fs_t* req, uv_file
file, uv_fs_cb cb);
int uv_fs_ftruncate(uv_loop_t* loop, uv_fs_t* req, uv_file
file, int64_t offset, uv_fs_cb cb);
int uv_fs_sendfile(uv_loop_t* loop, uv_fs_t* req, uv_file
out_fd, uv_file in_fd, int64_t in_offset, size_t length,
uv_fs_cb cb);
int uv_fs_access(uv_loop_t* loop, uv_fs_t* req, const char*
path, int mode, uv_fs_cb cb);
int uv_fs_chmod(uv_loop_t* loop, uv_fs_t* req, const char*
path, int mode, uv_fs_cb cb);
int uv_fs_utime(uv_loop_t* loop, uv_fs_t* req, const char*
path, double atime, double mtime, uv_fs_cb cb);
int uv_fs_futime(uv_loop_t* loop, uv_fs_t* req, uv_file file,
double atime, double mtime, uv_fs_cb cb);
int uv_fs_lutime(uv_loop_t* loop, uv_fs_t* req, const char*
path, double atime, double mtime, uv_fs_cb cb);
int uv_fs_lstat(uv_loop_t* loop, uv_fs_t* req, const char*
path, uv_fs_cb cb);
int uv_fs_link(uv_loop_t* loop, uv_fs_t* req, const char* path,
const char* new_path, uv_fs_cb cb);
int uv_fs_symlink(uv_loop_t* loop, uv_fs_t* req, const char*
path, const char* new_path, int flags, uv_fs_cb cb);
int uv_fs_readlink(uv_loop_t* loop, uv_fs_t* req, const char*
path, uv_fs_cb cb);
int uv_fs_realpath(uv_loop_t* loop, uv_fs_t* req, const char*
path, uv_fs_cb cb);
int uv_fs_fchmod(uv_loop_t* loop, uv_fs_t* req, uv_file file,
int mode, uv_fs_cb cb);
```

```
int uv_fs_chown(uv_loop_t* loop, uv_fs_t* req, const char*
path, uv_uid_t uid, uv_gid_t gid, uv_fs_cb cb);
int uv_fs_fchown(uv_loop_t* loop, uv_fs_t* req, uv_file file,
uv_uid_t uid, uv_gid_t gid, uv_fs_cb cb);
int uv_fs_lchown(uv_loop_t* loop, uv_fs_t* req, const char*
path, uv_uid_t uid, uv_gid_t gid, uv_fs_cb cb);
int uv_fs_statfs(uv_loop_t* loop, uv_fs_t* req, const char*
path, uv_fs_cb cb);
```

## Buffers and Streams

The basic I/O handle in libuv is the stream (`uv_stream_t`). TCP sockets, UDP sockets, and pipes for file I/O and IPC are all treated as stream subclasses.

Streams are initialized using custom functions for each subclass, then operated upon using

```
int uv_read_start(uv_stream_t*, uv_alloc_cb alloc_cb,
uv_read_cb read_cb);
int uv_read_stop(uv_stream_t*);
int uv_write(uv_write_t* req, uv_stream_t* handle,
const uv_buf_t bufs[], unsigned int nbufs,
uv_write_cb cb);
```

The stream based functions are simpler to use than the filesystem ones and libuv will automatically keep reading from a stream when `uv_read_start()` is called once, until `uv_read_stop()` is called.

The discrete unit of data is the buffer - `uv_buf_t`. This is simply a collection of a pointer to bytes (`uv_buf_t.base`) and the length (`uv_buf_t.len`). The `uv_buf_t` is lightweight and passed around by value. What does require management is the actual bytes, which have to be allocated and freed by the application.

### Error

## THIS PROGRAM DOES NOT ALWAYS WORK, NEED SOMETHING BETTER

To demonstrate streams we will need to use `uv_pipe_t`. This allows streaming local files [2]. Here is a simple tee utility using libuv. Doing all operations asynchronously shows the power of evented I/O. The two writes won't block each other, but we have to be careful to copy over the buffer data to ensure we don't free a buffer until it has been written.

The program is to be executed as:

```
./uvtee <output_file>
```

We start off opening pipes on the files we require. libuv pipes to a file are opened as bidirectional by default.

### uvtee/main.c - read on pipes

```
1     loop = uv_default_loop();
2
3     uv_pipe_init(loop, &stdin_pipe, 0);
4     uv_pipe_open(&stdin_pipe, 0);
5
6     uv_pipe_init(loop, &stdout_pipe, 0);
7     uv_pipe_open(&stdout_pipe, 1);
8
9     uv_fs_t file_req;
10    int fd = uv_fs_open(loop, &file_req, argv[1],
11    O_CREAT | O_RDWR, 0644, NULL);
12    uv_pipe_init(loop, &file_pipe, 0);
13    uv_pipe_open(&file_pipe, fd);
14
15    uv_read_start((uv_stream_t*)&stdin_pipe,
16    alloc_buffer, read_stdin);
17
18    uv_run(loop, UV_RUN_DEFAULT);
19    return 0;
20 }
```

The third argument of `uv_pipe_init()` should be set to 1 for IPC using named pipes. This is covered in [Processes](#). The `uv_pipe_open()` call associates the pipe with the file descriptor, in this case 0 (standard input).

We start monitoring `stdin`. The `alloc_buffer` callback is invoked as new buffers are required to hold incoming data. `read_stdin` will be called with these buffers.

### uvtee/main.c - reading buffers

```
    *buf = uv_buf_init((char*) malloc(suggested_size),
1   suggested_size);
2   }
3   void free_write_req(uv_write_t *req) {
4       if (nread < 0){
5           if (nread == UV_EOF){
6               // end of file
7               uv_close((uv_handle_t *)&stdin_pipe, NULL);
8               uv_close((uv_handle_t *)&stdout_pipe,
9   NULL);
10              uv_close((uv_handle_t *)&file_pipe, NULL);
11          }
12          } else if (nread > 0) {
13              write_data((uv_stream_t *)&stdout_pipe, nread,
14 *buf, on_stdout_write);
15              write_data((uv_stream_t *)&file_pipe, nread,
16 *buf, on_file_write);
17          }
18
19          // OK to free buffer as write_data copies it.
20          if (buf->base)
21              free(buf->base);
    }
```

The standard `malloc` is sufficient here, but you can use any memory allocation scheme. For example, `node.js` uses its own slab allocator which associates buffers with V8 objects.

The read callback `nread` parameter is less than 0 on any error. This error might be EOF, in which case we close all the streams, using the generic close function `uv_close()` which deals with the handle based on its internal type. Otherwise `nread` is a non-negative number and we can attempt to write that many bytes to the output streams. Finally remember that buffer allocation and deallocation is application responsibility, so we free the data.

The allocation callback may return a buffer with length zero if it fails to allocate memory. In this case, the read callback is invoked with error `UV_ENOBUFS`. `libuv` will continue to attempt to read the stream though, so you must explicitly call `uv_close()` if you want to stop when allocation fails.

The read callback may be called with `nread = 0`, indicating that at this point there is nothing to be read. Most applications will just ignore this.

## **uvtee/main.c - Write to pipe**

```
1     uv_write_t req;
2     uv_buf_t buf;
3 } write_req_t;
4
5 uv_loop_t *loop;
6     write_req_t *wr = (write_req_t*) req;
7     free(wr->buf.base);
8     free(wr);
9 }
10
11 void on_stdout_write(uv_write_t *req, int status) {
12     free_write_req(req);
13 }
```

```

14
15 void on_file_write(uv_write_t *req, int status) {
16     free_write_req(req);
17 }
18
19 void write_data(uv_stream_t *dest, size_t size,
20 uv_buf_t buf, uv_write_cb cb) {
21     write_req_t *req = (write_req_t*)
22     malloc(sizeof(write_req_t));
23     req->buf = uv_buf_init((char*) malloc(size), size);
24     memcpy(req->buf.base, buf.base, size);
25     uv_write((uv_write_t*) req, (uv_stream_t*)dest,
    &req->buf, 1, cb);
    }

```

write\_data() makes a copy of the buffer obtained from read. This buffer does not get passed through to the write callback triggered on write completion. To get around this we wrap a write request and a buffer in write\_req\_t and unwrap it in the callbacks. We make a copy so we can free the two buffers from the two calls to write\_data independently of each other. While acceptable for a demo program like this, you'll probably want smarter memory management, like reference counted buffers or a pool of buffers in any major application.

## Warning

If your program is meant to be used with other programs it may knowingly or unknowingly be writing to a pipe. This makes it susceptible to [aborting on receiving a SIGPIPE](#)

[[http://pod.tst.eu/http://cvs.schmorp.de/libev/ev.pod#The\\_special\\_problem\\_of\\_SIGPIPE](http://pod.tst.eu/http://cvs.schmorp.de/libev/ev.pod#The_special_problem_of_SIGPIPE)]. It is a good idea to insert:

```
signal(SIGPIPE, SIG_IGN)
```

in the initialization stages of your application.

# File change events

All modern operating systems provide APIs to put watches on individual files or directories and be informed when the files are modified. libuv wraps common file change notification libraries [1]. This is one of the more inconsistent parts of libuv. File change notification systems are themselves extremely varied across platforms so getting everything working everywhere is difficult. To demonstrate, I'm going to build a simple utility which runs a command whenever any of the watched files change:

```
./onchange <command> <file1> [file2] ...
```

## Note

Currently this example only works on OSX and Windows.

Refer to the [notes of uv\\_fs\\_event\\_start](#)

[[https://docs.libuv.org/en/v1.x/fs\\_event.html#c.uv\\_fs\\_event\\_start](https://docs.libuv.org/en/v1.x/fs_event.html#c.uv_fs_event_start)] function.

The file change notification is started using

```
uv_fs_event_init():
```

## onchange/main.c - The setup

```
1 int main(int argc, char **argv) {
2     if (argc <= 2) {
3         fprintf(stderr, "Usage: %s <command> <file1>
4 [file2 ...]\n", argv[0]);
5         return 1;
6     }
7
8     loop = uv_default_loop();
9     command = argv[1];
10
11     while (argc-- > 2) {
```



```

12     fprintf(stderr, "Adding watch on %s\n",
13     argv[argc]);
14     uv_fs_event_t *fs_event_req =
15     malloc(sizeof(uv_fs_event_t));
16     uv_fs_event_init(loop, fs_event_req);
17     // The recursive flag watches subdirectories
18     too.
19     uv_fs_event_start(fs_event_req, run_command,
    argv[argc], UV_FS_EVENT_RECURSIVE);
    }

    return uv_run(loop, UV_RUN_DEFAULT);
}

```

The third argument is the actual file or directory to monitor. The last argument, flags, can be:

```

/*
 * Flags to be passed to uv_fs_event_start().
 */
enum uv_fs_event_flags {
    UV_FS_EVENT_WATCH_ENTRY = 1,
    UV_FS_EVENT_STAT = 2,
    UV_FS_EVENT_RECURSIVE = 4
};

```

UV\_FS\_EVENT\_WATCH\_ENTRY and UV\_FS\_EVENT\_STAT don't do anything (yet). UV\_FS\_EVENT\_RECURSIVE will start watching subdirectories as well on supported platforms.

The callback will receive the following arguments:

1. uv\_fs\_event\_t \*handle - The handle. The path field of the handle is the file on which the watch was set.
2. const char \*filename - If a directory is being monitored, this is the file which was changed. Only non-null on Linux and Windows. May be null even on those platforms.
3. int events - one of UV\_RENAME or UV\_CHANGE, or a bitwise OR of both.

4. `int status` - If `status < 0`, there is an [libuv error](#).

In our example we simply print the arguments and run the command using `system()`.

## onchange/main.c - file change notification callback

```
void run_command(uv_fs_event_t *handle, const char
1 *filename, int events, int status) {
2     char path[1024];
3     size_t size = 1023;
4     // Does not handle error if path is longer than
5     1023.
6     uv_fs_event_getpath(handle, path, &size);
7     path[size] = '\0';
8
9     fprintf(stderr, "Change detected in %s: ", path);
10    if (events & UV_RENAME)
11        fprintf(stderr, "renamed");
12    if (events & UV_CHANGE)
13        fprintf(stderr, "changed");
14
15    fprintf(stderr, " %s\n", filename ? filename : "");
16    system(command);
}
```

---

[1] inotify on Linux, FSEvents on Darwin, kqueue on BSDs, ReadDirectoryChangesW on Windows, event ports on Solaris, unsupported on Cygwin

[2] see [Parent-child IPC](#)

# Networking

Networking in libuv is not much different from directly using the BSD socket interface, some things are easier, all are non-blocking, but the concepts stay the same. In addition libuv offers utility functions to abstract the annoying, repetitive and low-level tasks like setting up sockets using the BSD socket structures, DNS lookup, and tweaking various socket parameters.

The `uv_tcp_t` and `uv_udp_t` structures are used for network I/O.

## Note

The code samples in this chapter exist to show certain libuv APIs. They are not examples of good quality code. They leak memory and don't always close connections properly.

## TCP

TCP is a connection oriented, stream protocol and is therefore based on the libuv streams infrastructure.

## Server

Server sockets proceed by:

1. `uv_tcp_init` the TCP handle.
2. `uv_tcp_bind` it.
3. Call `uv_listen` on the handle to have a callback invoked whenever a new connection is established by a client.

4. Use `uv_accept` to accept the connection.
5. Use [stream operations](#) to communicate with the client.

Here is a simple echo server

### tcp-echo-server/main.c - The listen socket

```
    uv_close((uv_handle_t*) client, on_close);
1   }
2   }
3
4   int main() {
5       loop = uv_default_loop();
6
7       uv_tcp_t server;
8       uv_tcp_init(loop, &server);
9
10      uv_ip4_addr("0.0.0.0", DEFAULT_PORT, &addr);
11
12      uv_tcp_bind(&server, (const struct sockaddr*)&addr,
13      0);
14      int r = uv_listen((uv_stream_t*) &server,
15      DEFAULT_BACKLOG, on_new_connection);
16      if (r) {
17          fprintf(stderr, "Listen error %s\n",
18          uv_strerror(r));
19          return 1;
20      }
21      return uv_run(loop, UV_RUN_DEFAULT);
22  }
```

You can see the utility function `uv_ip4_addr` being used to convert from a human readable IP address, port pair to the `sockaddr_in` structure required by the BSD socket APIs. The reverse can be obtained using `uv_ip4_name`.

### Note

There are `uv_ip6_*` analogues for the ip4 functions.

Most of the setup functions are synchronous since they are CPU-bound. `uv_listen` is where we return to libuv's callback style. The second argument is the backlog queue - the maximum length of queued connections.

When a connection is initiated by clients, the callback is required to set up a handle for the client socket and associate the handle using `uv_accept`. In this case we also establish interest in reading from this stream.

### tcp-echo-server/main.c - Accepting the client

```
    free(buf->base);
1  }
2
3  void on_new_connection(uv_stream_t *server, int status)
4  {
5      if (status < 0) {
6          fprintf(stderr, "New connection error %s\n",
7  uv_strerror(status));
8          // error!
9          return;
10     }
11
12     uv_tcp_t *client = (uv_tcp_t*)
13     malloc(sizeof(uv_tcp_t));
14     uv_tcp_init(loop, client);
15     if (uv_accept(server, (uv_stream_t*) client) == 0)
16     {
17         uv_read_start((uv_stream_t*) client,
18         alloc_buffer, echo_read);
19     }
20 }
```

The remaining set of functions is very similar to the streams example and can be found in the code. Just remember to call `uv_close` when the socket isn't required. This can be

done even in the `uv_listen` callback if you are not interested in accepting the connection.

## Client

Where you do `bind/listen/accept` on the server, on the client side it's simply a matter of calling `uv_tcp_connect`. The same `uv_connect_cb` style callback of `uv_listen` is used by `uv_tcp_connect`. Try:

```
uv_tcp_t* socket = (uv_tcp_t*)malloc(sizeof(uv_tcp_t));
uv_tcp_init(loop, socket);

uv_connect_t* connect =
(uv_connect_t*)malloc(sizeof(uv_connect_t));

struct sockaddr_in dest;
uv_ip4_addr("127.0.0.1", 80, &dest);

uv_tcp_connect(connect, socket, (const struct sockaddr*)&dest,
on_connect);
```

where `on_connect` will be called after the connection is established. The callback receives the `uv_connect_t` struct, which has a member `.handle` pointing to the socket.

## UDP

The [User Datagram Protocol](https://en.wikipedia.org/wiki/User_Datagram_Protocol)

[[https://en.wikipedia.org/wiki/User\\_Datagram\\_Protocol](https://en.wikipedia.org/wiki/User_Datagram_Protocol)] offers connectionless, unreliable network communication. Hence `libuv` doesn't offer a stream. Instead `libuv` provides non-blocking UDP support via the `uv_udp_t` handle (for receiving) and `uv_udp_send_t` request (for sending) and related functions. That said, the actual API for reading/writing is very similar to normal stream reads. To look at how UDP can be used, the example shows the first stage of obtaining an

IP address from a [DHCP](https://tools.ietf.org/html/rfc2131) [https://tools.ietf.org/html/rfc2131] server – DHCP Discover.

## Note

You will have to run *udp-dhcp* as **root** since it uses well known port numbers below 1024.

## udp-dhcp/main.c - Setup and send UDP packets

```
1 uv_loop_t *loop;
2 uv_udp_t send_socket;
3 uv_udp_t recv_socket;
4
5
6
7 int main() {
8     loop = uv_default_loop();
9
10    uv_udp_init(loop, &recv_socket);
11    struct sockaddr_in recv_addr;
12    uv_ip4_addr("0.0.0.0", 68, &recv_addr);
13    uv_udp_bind(&recv_socket, (const struct sockaddr
14 *)&recv_addr, UV_UDP_REUSEADDR);
15    uv_udp_recv_start(&recv_socket, alloc_buffer,
16 on_read);
17
18    uv_udp_init(loop, &send_socket);
19    struct sockaddr_in broadcast_addr;
20    uv_ip4_addr("0.0.0.0", 0, &broadcast_addr);
21    uv_udp_bind(&send_socket, (const struct sockaddr
22 *)&broadcast_addr, 0);
23    uv_udp_set_broadcast(&send_socket, 1);
24
25    uv_udp_send_t send_req;
26    uv_buf_t discover_msg = make_discover_msg();
27
28    struct sockaddr_in send_addr;
29    uv_ip4_addr("255.255.255.255", 67, &send_addr);
30    uv_udp_send(&send_req, &send_socket, &discover_msg,
```

```
1, (const struct sockaddr *)&send_addr, on_send);  
  
    return uv_run(loop, UV_RUN_DEFAULT);  
}
```

## Note

The IP address 0.0.0.0 is used to bind to all interfaces. The IP address 255.255.255.255 is a broadcast address meaning that packets will be sent to all interfaces on the subnet. port 0 means that the OS randomly assigns a port.

First we setup the receiving socket to bind on all interfaces on port 68 (DHCP client) and start a read on it. This will read back responses from any DHCP server that replies. We use the UV\_UDP\_REUSEADDR flag to play nice with any other system DHCP clients that are running on this computer on the same port. Then we setup a similar send socket and use uv\_udp\_send to send a *broadcast message* on port 67 (DHCP server).

It is **necessary** to set the broadcast flag, otherwise you will get an EACCES error [1]. The exact message being sent is not relevant to this book and you can study the code if you are interested. As usual the read and write callbacks will receive a status code of < 0 if something went wrong.

Since UDP sockets are not connected to a particular peer, the read callback receives an extra parameter about the sender of the packet.

nread may be zero if there is no more data to be read. If addr is NULL, it indicates there is nothing to read (the callback shouldn't do anything), if not NULL, it indicates that an empty datagram was received from the host at addr. The flags parameter may be UV\_UDP\_PARTIAL if the buffer provided



by your allocator was not large enough to hold the data. *In this case the OS will discard the data that could not fit* (That's UDP for you!).

## udp-dhcp/main.c - Reading packets

```
void on_read(uv_udp_t *req, ssize_t nread, const
uv_buf_t *buf, const struct sockaddr *addr, unsigned
flags) {
1   if (nread < 0) {
2       fprintf(stderr, "Read error %s\n",
3   uv_err_name(nread));
4       uv_close((uv_handle_t*) req, NULL);
5       free(buf->base);
6       return;
7   }
8
9   char sender[17] = { 0 };
10  uv_ip4_name((const struct sockaddr_in*) addr,
11  sender, 16);
12  fprintf(stderr, "Recv from %s\n", sender);
13
14  // ... DHCP specific code
15  unsigned int *as_integer = (unsigned int*)buf-
16  >base;
17  unsigned int ipbin = ntohl(as_integer[4]);
18  unsigned char ip[4] = {0};
19  int i;
20  for (i = 0; i < 4; i++)
21      ip[i] = (ipbin >> i*8) & 0xff;
22  fprintf(stderr, "Offered IP %d.%d.%d.%d\n", ip[3],
23  ip[2], ip[1], ip[0]);
24
25  free(buf->base);
26  uv_udp_recv_stop(req);
}
```

## UDP Options

### Time-to-live

The TTL of packets sent on the socket can be changed using `uv_udp_set_ttl`.

## IPv6 stack only

IPv6 sockets can be used for both IPv4 and IPv6 communication. If you want to restrict the socket to IPv6 only, pass the `UV_UDP_IPV6ONLY` flag to `uv_udp_bind` [2].

## Multicast

A socket can (un)subscribe to a multicast group using:

where `membership` is `UV_JOIN_GROUP` or `UV_LEAVE_GROUP`.

The concepts of multicasting are nicely explained in [this guide](https://www.tldp.org/HOWTO/Multicast-HOWTO-2.html) [https://www.tldp.org/HOWTO/Multicast-HOWTO-2.html].

Local loopback of multicast packets is enabled by default [3], use `uv_udp_set_multicast_loop` to switch it off.

The packet time-to-live for multicast packets can be changed using `uv_udp_set_multicast_ttl`.

## Querying DNS

libuv provides asynchronous DNS resolution. For this it provides its own `getaddrinfo` replacement [4]. In the callback you can perform normal socket operations on the retrieved addresses. Let's connect to Libera.chat to see an example of DNS resolution.

### dns/main.c

```

2  int main() {
3      loop = uv_default_loop();
4
5      struct addrinfo hints;
6      hints.ai_family = PF_INET;
7      hints.ai_socktype = SOCK_STREAM;
8      hints.ai_protocol = IPPROTO_TCP;
9      hints.ai_flags = 0;
10
11     uv_getaddrinfo_t resolver;
12     fprintf(stderr, "irc.libera.chat is... ");
13     int r = uv_getaddrinfo(loop, &resolver,
14 on_resolved, "irc.libera.chat", "6667", &hints);
15
16     if (r) {
17         fprintf(stderr, "getaddrinfo call error %s\n",
18 uv_err_name(r));
19         return 1;
20     }
21     return uv_run(loop, UV_RUN_DEFAULT);
22 }

```

If `uv_getaddrinfo` returns non-zero, something went wrong in the setup and your callback won't be invoked at all. All arguments can be freed immediately after `uv_getaddrinfo` returns. The *hostname*, *servname* and *hints* structures are documented in [the getaddrinfo man page](https://man7.org/linux/man-pages/man3/getaddrinfo.3.html)

[<https://man7.org/linux/man-pages/man3/getaddrinfo.3.html>]. The callback can be NULL in which case the function will run synchronously.

In the resolver callback, you can pick any IP from the linked list of `struct addrinfo(s)`. This also demonstrates `uv_tcp_connect`. It is necessary to call `uv_freeaddrinfo` in the callback.

## dns/main.c

```

2 void on_resolved(uv_getaddrinfo_t *resolver, int
3 status, struct addrinfo *res) {
4     if (status < 0) {
5         fprintf(stderr, "getaddrinfo callback error
6 %s\n", uv_err_name(status));
7         return;
8     }
9
10    char addr[17] = {'\0'};
11    uv_ip4_name((struct sockaddr_in*) res->ai_addr,
12 addr, 16);
13    fprintf(stderr, "%s\n", addr);
14
15    uv_connect_t *connect_req = (uv_connect_t*)
16 malloc(sizeof(uv_connect_t));
17    uv_tcp_t *socket = (uv_tcp_t*)
18 malloc(sizeof(uv_tcp_t));
19    uv_tcp_init(loop, socket);

    uv_tcp_connect(connect_req, socket, (const struct
sockaddr*) res->ai_addr, on_connect);

    uv_freeaddrinfo(res);
}

```

libuv also provides the inverse [uv\\_getnameinfo](#)

[[http://docs.libuv.org/en/v1.x/dns.html#c.uv\\_getnameinfo](http://docs.libuv.org/en/v1.x/dns.html#c.uv_getnameinfo)].

## Network interfaces

Information about the system's network interfaces can be obtained through libuv using `uv_interface_addresses`. This simple program just prints out all the interface details so you get an idea of the fields that are available. This is useful to allow your service to bind to IP addresses when it starts.

### interfaces/main.c

```

1 #include <stdio.h>

```

```

2  #include <uv.h>
3
4  int main() {
5      char buf[512];
6      uv_interface_address_t *info;
7      int count, i;
8
9      uv_interface_addresses(&info, &count);
10     i = count;
11
12     printf("Number of interfaces: %d\n", count);
13     while (i--) {
14         uv_interface_address_t interface_a = info[i];
15
16         printf("Name: %s\n", interface_a.name);
17         printf("Internal? %s\n",
18 interface_a.is_internal ? "Yes" : "No");
19
20         if (interface_a.address.address4.sin_family ==
21 AF_INET) {
22             uv_ip4_name(&interface_a.address.address4,
23 buf, sizeof(buf));
24             printf("IPv4 address: %s\n", buf);
25         }
26         else if
27 (interface_a.address.address4.sin_family == AF_INET6) {
28             uv_ip6_name(&interface_a.address.address6,
29 buf, sizeof(buf));
30             printf("IPv6 address: %s\n", buf);
31         }
32
33         printf("\n");
34
35         uv_free_interface_addresses(info, count);
36         return 0;
37     }
38 }

```

`is_internal` is true for loopback interfaces. Note that if a physical interface has multiple IPv4/IPv6 addresses, the name will be reported multiple times, with each address being reported once.

---

- [1] <https://beej.us/guide/bgnet/html/#broadcast-packetshello-world>
- [2] on Windows only supported on Windows Vista and later.
- [3] <https://www.tldp.org/HOWTO/Multicast-HOWTO-6.html#ss6.1>
- [4] libuv use the system `getaddrinfo` in the libuv threadpool. libuv v0.8.0 and earlier also included [c-ares](https://c-ares.haxx.se) [https://c-ares.haxx.se] as an alternative, but this has been removed in v0.9.0.

# Threads

Wait a minute? Why are we on threads? Aren't event loops supposed to be **the way** to do *web-scale programming*? Well... no. Threads are still the medium in which processors do their jobs. Threads are therefore mighty useful sometimes, even though you might have to wade through various synchronization primitives.

Threads are used internally to fake the asynchronous nature of all of the system calls. libuv also uses threads to allow you, the application, to perform a task asynchronously that is actually blocking, by spawning a thread and collecting the result when it is done.

Today there are two predominant thread libraries: the Windows threads implementation and POSIX's [pthreads\(7\)](https://man7.org/linux/man-pages/man7/pthreads.7.html). [https://man7.org/linux/man-pages/man7/pthreads.7.html]. libuv's thread API is analogous to the pthreads API and often has similar semantics.

A notable aspect of libuv's thread facilities is that it is a self contained section within libuv. Whereas other features intimately depend on the event loop and callback principles, threads are complete agnostic, they block as required, signal errors directly via return values, and, as shown in the [first example](#), don't even require a running event loop.

libuv's thread API is also very limited since the semantics and syntax of threads are different on all platforms, with different levels of completeness.

This chapter makes the following assumption: **There is only one event loop, running in one thread (the main**

**thread**). No other thread interacts with the event loop (except using `uv_async_send`).

## Core thread operations

There isn't much here, you just start a thread using `uv_thread_create()` and wait for it to close using `uv_thread_join()`.

### thread-create/main.c

```
1   int tracklen = 10;
2   uv_thread_t hare_id;
3   uv_thread_t tortoise_id;
4   uv_thread_create(&hare_id, hare, &tracklen);
5   uv_thread_create(&tortoise_id, tortoise,
6   &tracklen);
7
8   uv_thread_join(&hare_id);
9   uv_thread_join(&tortoise_id);
10  return 0;
}
```

### Tip

`uv_thread_t` is just an alias for `pthread_t` on Unix, but this is an implementation detail, avoid depending on it to always be true.

The second parameter is the function which will serve as the entry point for the thread, the last parameter is a `void *` argument which can be used to pass custom parameters to the thread. The function `hare` will now run in a separate thread, scheduled pre-emptively by the operating system:



## thread-create/main.c

```
1     int tracklen = *((int *) arg);
2     while (tracklen) {
3         tracklen--;
4         uv_sleep(1000);
5         fprintf(stderr, "Hare ran another step\n");
6     }
7     fprintf(stderr, "Hare done running!\n");
8 }
9
```

Unlike `pthread_join()` which allows the target thread to pass back a value to the calling thread using a second parameter, `uv_thread_join()` does not. To send values use [Inter-thread communication](#).

## Synchronization Primitives

This section is purposely spartan. This book is not about threads, so I only catalogue any surprises in the libuv APIs here. For the rest you can look at the [pthreads\(7\)](#).

[<https://man7.org/linux/man-pages/man7/pthreads.7.html>] man pages.

### Mutexes

The mutex functions are a **direct** map to the pthread equivalents.

### libuv mutex functions

```
int uv_mutex_init(uv_mutex_t* handle);
int uv_mutex_init_recursive(uv_mutex_t* handle);
void uv_mutex_destroy(uv_mutex_t* handle);
void uv_mutex_lock(uv_mutex_t* handle);
int uv_mutex_trylock(uv_mutex_t* handle);
void uv_mutex_unlock(uv_mutex_t* handle);
```

The `uv_mutex_init()`, `uv_mutex_init_recursive()` and `uv_mutex_trylock()` functions will return 0 on success, and an error code otherwise.

If *libuv* has been compiled with debugging enabled, `uv_mutex_destroy()`, `uv_mutex_lock()` and `uv_mutex_unlock()` will `abort()` on error. Similarly `uv_mutex_trylock()` will abort if the error is anything *other than* `EAGAIN` or `EBUSY`.

Recursive mutexes are supported, but you should not rely on them. Also, they should not be used with `uv_cond_t` variables.

The default BSD mutex implementation will raise an error if a thread which has locked a mutex attempts to lock it again. For example, a construct like:

```
uv_mutex_init(a_mutex);
uv_mutex_lock(a_mutex);
uv_thread_create(thread_id, entry, (void *)a_mutex);
uv_mutex_lock(a_mutex);
// more things here
```

can be used to wait until another thread initializes some stuff and then unlocks `a_mutex` but will lead to your program crashing if in debug mode, or return an error in the second call to `uv_mutex_lock()`.

## Note

Mutexes on Windows are always recursive.

## Locks

Read-write locks are a more granular access mechanism. Two readers can access shared memory at the same time. A

writer may not acquire the lock when it is held by a reader. A reader or writer may not acquire a lock when a writer is holding it. Read-write locks are frequently used in databases. Here is a toy example.

## locks/main.c - simple rwlocks

```
1  #include <stdio.h>
2  #include <uv.h>
3
4  uv_barrier_t blocker;
5  uv_rwlock_t numlock;
6  int shared_num;
7
8  void reader(void *n)
9  {
10     int num = *(int *)n;
11     int i;
12     for (i = 0; i < 20; i++) {
13         uv_rwlock_rdlock(&numlock);
14         printf("Reader %d: acquired lock\n", num);
15         printf("Reader %d: shared num = %d\n", num,
16 shared_num);
17         uv_rwlock_runlock(&numlock);
18         printf("Reader %d: released lock\n", num);
19     }
20     uv_barrier_wait(&blocker);
21 }
22
23 void writer(void *n)
24 {
25     int num = *(int *)n;
26     int i;
27     for (i = 0; i < 20; i++) {
28         uv_rwlock_wrlock(&numlock);
29         printf("Writer %d: acquired lock\n", num);
30         shared_num++;
31         printf("Writer %d: incremented shared num =
32 %d\n", num, shared_num);
33         uv_rwlock_wrunlock(&numlock);
34         printf("Writer %d: released lock\n", num);
35     }
}
```

```

36     uv_barrier_wait(&blocker);
37 }
38
39 int main()
40 {
41     uv_barrier_init(&blocker, 4);
42
43     shared_num = 0;
44     uv_rwlock_init(&numlock);
45
46     uv_thread_t threads[3];
47
48     int thread_nums[] = {1, 2, 1};
49     uv_thread_create(&threads[0], reader,
50 &thread_nums[0]);
51     uv_thread_create(&threads[1], reader,
52 &thread_nums[1]);
53
54     uv_thread_create(&threads[2], writer,
55 &thread_nums[2]);
56
57     uv_barrier_wait(&blocker);
58     uv_barrier_destroy(&blocker);
59
60     uv_rwlock_destroy(&numlock);
61     return 0;
62 }

```

Run this and observe how the readers will sometimes overlap. In case of multiple writers, schedulers will usually give them higher priority, so if you add two writers, you'll see that both writers tend to finish first before the readers get a chance again.

We also use barriers in the above example so that the main thread can wait for all readers and writers to indicate they have ended.

## Others

libuv also supports [semaphores](#)

[[https://en.wikipedia.org/wiki/Semaphore\\_\(programming\)](https://en.wikipedia.org/wiki/Semaphore_(programming))], [condition variables](#)

[[https://en.wikipedia.org/wiki/Monitor\\_\(synchronization\)#Condition\\_variables\\_2](https://en.wikipedia.org/wiki/Monitor_(synchronization)#Condition_variables_2)] and [barriers](#) [[https://en.wikipedia.org/wiki/Barrier\\_\(computer\\_science\)](https://en.wikipedia.org/wiki/Barrier_(computer_science))] with APIs very similar to their pthread counterparts.

In addition, libuv provides a convenience function `uv_once()`. Multiple threads can attempt to call `uv_once()` with a given guard and a function pointer, **only the first one will win, the function will be called once and only once:**

```
/* Initialize guard */
static uv_once_t once_only = UV_ONCE_INIT;

int i = 0;

void increment() {
    i++;
}

void thread1() {
    /* ... work */
    uv_once(once_only, increment);
}

void thread2() {
    /* ... work */
    uv_once(once_only, increment);
}

int main() {
    /* ... spawn threads */
}
```

After all threads are done, `i == 1`.

libuv v0.11.11 onwards also added a `uv_key_t` struct and [api](#) [<http://docs.libuv.org/en/v1.x/threading.html#thread-local-storage>] for thread-local storage.

# libuv work queue

`uv_queue_work()` is a convenience function that allows an application to run a task in a separate thread, and have a callback that is triggered when the task is done. A seemingly simple function, what makes `uv_queue_work()` tempting is that it allows potentially any third-party libraries to be used with the event-loop paradigm. When you use event loops, it is *imperative to make sure that no function which runs periodically in the loop thread blocks when performing I/O or is a serious CPU hog*, because this means that the loop slows down and events are not being handled at full capacity.

However, a lot of existing code out there features blocking functions (for example a routine which performs I/O under the hood) to be used with threads if you want responsiveness (the classic 'one thread per client' server model), and getting them to play with an event loop library generally involves rolling your own system of running the task in a separate thread. libuv just provides a convenient abstraction for this.

Here is a simple example inspired by [node.js is cancer](http://widgetsandshit.com/teddziuba/2011/10/node-js-is-cancer.html) [http://widgetsandshit.com/teddziuba/2011/10/node-js-is-cancer.html]. We are going to calculate fibonacci numbers, sleeping a bit along the way, but run it in a separate thread so that the blocking and CPU bound task does not prevent the event loop from performing other activities.

## queue-work/main.c - lazy fibonacci

```
1 void fib(uv_work_t *req) {  
2     int n = *(int *) req->data;  
3     if (random() % 2)
```

```

4     sleep(1);
5     else
6         sleep(3);
7     long fib = fib_(n);
8     fprintf(stderr, "%dth fibonacci is %lu\n", n, fib);
9 }
10
11 void after_fib(uv_work_t *req, int status) {
12     fprintf(stderr, "Done calculating %dth
13 fibonacci\n", *(int *) req->data);
14 }

```

The actual task function is simple, nothing to show that it is going to be run in a separate thread. The `uv_work_t` structure is the clue. You can pass arbitrary data through it using the `void* data` field and use it to communicate to and from the thread. But be sure you are using proper locks if you are changing things while both threads may be running.

The trigger is `uv_queue_work`:

### queue-work/main.c

```

1 int main() {
2     loop = uv_default_loop();
3
4     int data[FIB_UNTIL];
5     uv_work_t req[FIB_UNTIL];
6     int i;
7     for (i = 0; i < FIB_UNTIL; i++) {
8         data[i] = i;
9         req[i].data = (void *) &data[i];
10        uv_queue_work(loop, &req[i], fib, after_fib);
11    }
12
13    return uv_run(loop, UV_RUN_DEFAULT);
14 }

```

The thread function will be launched in a separate thread, passed the `uv_work_t` structure and once the function returns, the *after* function will be called on the thread the event loop is running in. It will be passed the same structure.

For writing wrappers to blocking libraries, a common [pattern](#) is to use a baton to exchange data.

Since libuv version *0.9.4* an additional function, `uv_cancel()`, is available. This allows you to cancel tasks on the libuv work queue. Only tasks that *are yet to be started* can be cancelled. If a task has *already started executing*, or it has *finished executing*, `uv_cancel()` **will fail**.

`uv_cancel()` is useful to cleanup pending tasks if the user requests termination. For example, a music player may queue up multiple directories to be scanned for audio files. If the user terminates the program, it should quit quickly and not wait until all pending requests are run.

Let's modify the fibonacci example to demonstrate `uv_cancel()`. We first set up a signal handler for termination.

### queue-cancel/main.c

```
1  int main() {
2      loop = uv_default_loop();
3
4      int data[FIB_UNTIL];
5      int i;
6      for (i = 0; i < FIB_UNTIL; i++) {
7          data[i] = i;
8          fib_reqs[i].data = (void *) &data[i];
9          uv_queue_work(loop, &fib_reqs[i], fib,
10         after_fib);
11     }
12
13     uv_signal_t sig;
```



```

14     uv_signal_init(loop, &sig);
15     uv_signal_start(&sig, signal_handler, SIGINT);
16
17     return uv_run(loop, UV_RUN_DEFAULT);
}

```

When the user triggers the signal by pressing Ctrl+C we send `uv_cancel()` to all the workers. `uv_cancel()` will return 0 for those that are already executing or finished.

### queue-cancel/main.c

```

1 void signal_handler(uv_signal_t *req, int signum)
2 {
3     printf("Signal received!\n");
4     int i;
5     for (i = 0; i < FIB_UNTIL; i++) {
6         uv_cancel((uv_req_t*) &fib_reqs[i]);
7     }
8     uv_signal_stop(req);
9 }

```

For tasks that do get cancelled successfully, the *after* function is called with `status` set to `UV_ECANCELED`.

### queue-cancel/main.c

```

1 void after_fib(uv_work_t *req, int status) {
2     if (status == UV_ECANCELED)
3         fprintf(stderr, "Calculation of %d
4 cancelled.\n", *(int *) req->data);
}

```

`uv_cancel()` can also be used with `uv_fs_t` and `uv_getaddrinfo_t` requests. For the filesystem family of functions, `uv_fs_t.errno` will be set to `UV_ECANCELED`.

## Tip

A well designed program would have a way to terminate long running workers that have already started executing. Such a worker could periodically check for a variable that only the main process sets to signal termination.

## Inter-thread communication

Sometimes you want various threads to actually send each other messages *while* they are running. For example you might be running some long duration task in a separate thread (perhaps using `uv_queue_work`) but want to notify progress to the main thread. This is a simple example of having a download manager informing the user of the status of running downloads.

### progress/main.c

```
1 uv_loop_t *loop;
2 uv_async_t async;
3
4 int main() {
5     loop = uv_default_loop();
6
7     uv_work_t req;
8     int size = 10240;
9     req.data = (void*) &size;
10
11     uv_async_init(loop, &async, print_progress);
12     uv_queue_work(loop, &req, fake_download, after);
13
14     return uv_run(loop, UV_RUN_DEFAULT);
15 }
```

The async thread communication works *on loops* so although any thread can be the message sender, only threads with libuv loops can be receivers (or rather the loop is the receiver). libuv will invoke the callback (`print_progress`) with the async watcher whenever it receives a message.

## Warning

It is important to realize that since the message send is *async*, the callback may be invoked immediately after `uv_async_send` is called in another thread, or it may be invoked after some time. libuv may also combine multiple calls to `uv_async_send` and invoke your callback only once. The only guarantee that libuv makes is – The callback function is called *at least once* after the call to `uv_async_send`. If you have no pending calls to `uv_async_send`, the callback won't be called. If you make two or more calls, and libuv hasn't had a chance to run the callback yet, it *may* invoke your callback *only once* for the multiple invocations of `uv_async_send`. Your callback will never be called twice for just one event.

## progress/main.c

```
1  double percentage;
2
3  void fake_download(uv_work_t *req) {
4      int size = *((int*) req->data);
5      int downloaded = 0;
6      while (downloaded < size) {
7          percentage = downloaded*100.0/size;
8          async.data = (void*) &percentage;
9          uv_async_send(&async);
10
11         sleep(1);
12         downloaded += (200+random())%1000; // can only
```

```

13  download max 1000bytes/sec,
14                                     // but at
15  least a 200;
    }
}

```

In the download function, we modify the progress indicator and queue the message for delivery with `uv_async_send`. Remember: `uv_async_send` is also non-blocking and will return immediately.

### progress/main.c

```

1  void print_progress(uv_async_t *handle) {
2      double percentage = *((double*) handle->data);
3      fprintf(stderr, "Downloaded %.2f%%\n", percentage);
4  }

```

The callback is a standard libuv pattern, extracting the data from the watcher.

Finally it is important to remember to clean up the watcher.

### progress/main.c

```

1  void after(uv_work_t *req, int status) {
2      fprintf(stderr, "Download complete\n");
3      uv_close((uv_handle_t*) &async, NULL);
4  }

```

After this example, which showed the abuse of the data field, [bnoordhuis](https://github.com/bnoordhuis) [https://github.com/bnoordhuis] pointed out that using the data field is not thread safe, and `uv_async_send()` is actually only meant to wake up the event loop. Use a mutex

or `rwlock` to ensure accesses are performed in the right order.

## Note

mutexes and `rwlocks` **DO NOT** work inside a signal handler, whereas `uv_async_send` does.

One use case where `uv_async_send` is required is when interoperating with libraries that require thread affinity for their functionality. For example in `node.js`, a `v8` engine instance, contexts and its objects are bound to the thread that the `v8` instance was started in. Interacting with `v8` data structures from another thread can lead to undefined results. Now consider some `node.js` module which binds a third party library. It may go something like this:

1. In `node`, the third party library is set up with a JavaScript callback to be invoked for more information:

```
var lib = require('lib');
lib.on_progress(function() {
  console.log("Progress");
});

lib.do();

// do other stuff
```

2. `lib.do` is supposed to be non-blocking but the third party lib is blocking, so the binding uses `uv_queue_work`.
3. The actual work being done in a separate thread wants to invoke the progress callback, but cannot directly call into `v8` to interact with JavaScript. So it uses `uv_async_send`.

4. The async callback, invoked in the main loop thread, which is the v8 thread, then interacts with v8 to invoke the JavaScript callback.
-

# Processes

libuv offers considerable child process management, abstracting the platform differences and allowing communication with the child process using streams or named pipes.

A common idiom in Unix is for every process to do one thing and do it well. In such a case, a process often uses multiple child processes to achieve tasks (similar to using pipes in shells). A multi-process model with messages may also be easier to reason about compared to one with threads and shared memory.

A common refrain against event-based programs is that they cannot take advantage of multiple cores in modern computers. In a multi-threaded program the kernel can perform scheduling and assign different threads to different cores, improving performance. But an event loop has only one thread. The workaround can be to launch multiple processes instead, with each process running an event loop, and each process getting assigned to a separate CPU core.

## Spawning child processes

The simplest case is when you simply want to launch a process and know when it exits. This is achieved using `uv_spawn`.

### `spawn/main.c`

```
1 uv_loop_t *loop;  
2 uv_process_t child_req;
```

```

3 uv_process_options_t options;
4 int main() {
5     loop = uv_default_loop();
6
7     char* args[3];
8     args[0] = "mkdir";
9     args[1] = "test-dir";
10    args[2] = NULL;
11
12    options.exit_cb = on_exit;
13    options.file = "mkdir";
14    options.args = args;
15
16    int r;
17    if ((r = uv_spawn(loop, &child_req, &options))) {
18        fprintf(stderr, "%s\n", uv_strerror(r));
19        return 1;
20    } else {
21        fprintf(stderr, "Launched process with ID
22 %d\n", child_req.pid);
23    }
24
25    return uv_run(loop, UV_RUN_DEFAULT);
}

```

## Note

options is implicitly initialized with zeros since it is a global variable. If you change options to a local variable, remember to initialize it to null out all unused fields:

```
uv_process_options_t options = {0};
```

The `uv_process_t` struct only acts as the handle, all options are set via `uv_process_options_t`. To simply launch a process, you need to set only the `file` and `args` fields. `file` is the program to execute. Since `uv_spawn` uses [execvp\(3\)](https://man7.org/linux/man-pages/man3/execvp.3.html) internally, there is no need to supply the full path. Finally as per underlying conventions, **the arguments array has to be one larger**



**than the number of arguments, with the last element being NULL.**

After the call to `uv_spawn`, `uv_process_t.pid` will contain the process ID of the child process.

The exit callback will be invoked with the *exit status* and the type of *signal* which caused the exit.

### **spawn/main.c**

```
1 void on_exit(uv_process_t *req, int64_t exit_status, int
2 term_signal) {
3     fprintf(stderr, "Process exited with status %"
4     PRIu64 ", signal %d\n", exit_status, term_signal);
5     uv_close((uv_handle_t*) req, NULL);
```

It is **required** to close the process watcher after the process exits.

## **Changing process parameters**

Before the child process is launched you can control the execution environment using fields in `uv_process_options_t`.

### **Change execution directory**

Set `uv_process_options_t.cwd` to the corresponding directory.

### **Set environment variables**

`uv_process_options_t.env` is a null-terminated array of strings, each of the form `VAR=VALUE` used to set up the environment

variables for the process. Set this to `NULL` to inherit the environment from the parent (this) process.

## Option flags

Setting `uv_process_options_t.flags` to a bitwise OR of the following flags, modifies the child process behaviour:

- `UV_PROCESS_SETUID` - sets the child's execution user ID to `uv_process_options_t.uid`.
- `UV_PROCESS_SETGID` - sets the child's execution group ID to `uv_process_options_t.gid`.

Changing the UID/GID is only supported on Unix, `uv_spawn` will fail on Windows with `UV_ENOTSUP`.

- `UV_PROCESS_WINDOWS_VERBATIM_ARGUMENTS` - No quoting or escaping of `uv_process_options_t.args` is done on Windows. Ignored on Unix.
- `UV_PROCESS_DETACHED` - Starts the child process in a new session, which will keep running after the parent process exits. See example below.

## Detaching processes

Passing the flag `UV_PROCESS_DETACHED` can be used to launch daemons, or child processes which are independent of the parent so that the parent exiting does not affect it.

### **detach/main.c**

```
1 int main() {
2     loop = uv_default_loop();
3
4     char* args[3];
5     args[0] = "sleep";
```

```

6     args[1] = "100";
7     args[2] = NULL;
8
9     options.exit_cb = NULL;
10    options.file = "sleep";
11    options.args = args;
12    options.flags = UV_PROCESS_DETACHED;
13
14    int r;
15    if ((r = uv_spawn(loop, &child_req, &options))) {
16        fprintf(stderr, "%s\n", uv_strerror(r));
17        return 1;
18    }
19    fprintf(stderr, "Launched sleep with PID %d\n",
20    child_req.pid);
21    uv_unref((uv_handle_t*) &child_req);
22
23    return uv_run(loop, UV_RUN_DEFAULT);

```

Just remember that the handle is still monitoring the child, so your program won't exit. Use `uv_unref()` if you want to be more *fire-and-forget*.

## Sending signals to processes

libuv wraps the standard `kill(2)` system call on Unix and implements one with similar semantics on Windows, with *one caveat*: all of `SIGTERM`, `SIGINT` and `SIGKILL`, lead to termination of the process. The signature of `uv_kill` is:

```
uv_err_t uv_kill(int pid, int signum);
```

For processes started using libuv, you may use `uv_process_kill` instead, which accepts the `uv_process_t` watcher as the first argument, rather than the pid. In this case, **remember to call** `uv_close` on the watcher.

## Signals

libuv provides wrappers around Unix signals with [some Windows support](http://docs.libuv.org/en/v1.x/signal.html#signal) [http://docs.libuv.org/en/v1.x/signal.html#signal] as well.

Use `uv_signal_init()` to initialize a handle and associate it with a loop. To listen for particular signals on that handler, use `uv_signal_start()` with the handler function. Each handler can only be associated with one signal number, with subsequent calls to `uv_signal_start()` overwriting earlier associations. Use `uv_signal_stop()` to stop watching. Here is a small example demonstrating the various possibilities:

### signal/main.c

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <unistd.h>
4  #include <uv.h>
5
6  uv_loop_t* create_loop()
7  {
8      uv_loop_t *loop = malloc(sizeof(uv_loop_t));
9      if (loop) {
10         uv_loop_init(loop);
11     }
12     return loop;
13 }
14
15 void signal_handler(uv_signal_t *handle, int signum)
16 {
17     printf("Signal received: %d\n", signum);
18     uv_signal_stop(handle);
19 }
20
21 // two signal handlers in one loop
22 void thread1_worker(void *userp)
23 {
24     uv_loop_t *loop1 = create_loop();
25
26     uv_signal_t sig1a, sig1b;
```

```

27     uv_signal_init(loop1, &sig1a);
28     uv_signal_start(&sig1a, signal_handler, SIGUSR1);
29
30     uv_signal_init(loop1, &sig1b);
31     uv_signal_start(&sig1b, signal_handler, SIGUSR1);
32
33     uv_run(loop1, UV_RUN_DEFAULT);
34 }
35
36 // two signal handlers, each in its own loop
37 void thread2_worker(void *userp)
38 {
39     uv_loop_t *loop2 = create_loop();
40     uv_loop_t *loop3 = create_loop();
41
42     uv_signal_t sig2;
43     uv_signal_init(loop2, &sig2);
44     uv_signal_start(&sig2, signal_handler, SIGUSR1);
45
46     uv_signal_t sig3;
47     uv_signal_init(loop3, &sig3);
48     uv_signal_start(&sig3, signal_handler, SIGUSR1);
49
50     while (uv_run(loop2, UV_RUN_NOWAIT) ||
51 uv_run(loop3, UV_RUN_NOWAIT)) {
52     }
53 }
54
55 int main()
56 {
57     printf("PID %d\n", getpid());
58
59     uv_thread_t thread1, thread2;
60
61     uv_thread_create(&thread1, thread1_worker, 0);
62     uv_thread_create(&thread2, thread2_worker, 0);
63
64     uv_thread_join(&thread1);
65     uv_thread_join(&thread2);
66     return 0;
}

```

## Note

`uv_run(loop, UV_RUN_NOWAIT)` is similar to `uv_run(loop, UV_RUN_ONCE)` in that it will process only one event. `UV_RUN_ONCE` blocks if there are no pending events, while `UV_RUN_NOWAIT` will return immediately. We use `NOWAIT` so that one of the loops isn't starved because the other one has no pending activity.

Send `SIGUSR1` to the process, and you'll find the handler being invoked 4 times, one for each `uv_signal_t`. The handler just stops each handle, so that the program exits. This sort of dispatch to all handlers is very useful. A server using multiple event loops could ensure that all data was safely saved before termination, simply by every loop adding a watcher for `SIGINT`.

## Child Process I/O

A normal, newly spawned process has its own set of file descriptors, with 0, 1 and 2 being `stdin`, `stdout` and `stderr` respectively. Sometimes you may want to share file descriptors with the child. For example, perhaps your application launches a sub-command and you want any errors to go in the log file, but ignore `stdout`. For this you'd like to have `stderr` of the child be the same as the `stderr` of the parent. In this case, `libuv` supports *inheriting* file descriptors. In this sample, we invoke the test program, which is:

### **proc-streams/test.c**

```
#include <stdio.h>

int main()
{
    fprintf(stderr, "This is stderr\n");
}
```

```
    printf("This is stdout\n");
    return 0;
}
```

The actual program `proc-streams` runs this while sharing only `stderr`. The file descriptors of the child process are set using the `stdio` field in `uv_process_options_t`. First set the `stdio_count` field to the number of file descriptors being set.

`uv_process_options_t.stdio` is an array of `uv_stdio_container_t`, which is:

```
typedef struct uv_stdio_container_s {
    uv_stdio_flags flags;

    union {
        uv_stream_t* stream;
        int fd;
    } data;
} uv_stdio_container_t;
```

where `flags` can have several values. Use `UV_IGNORE` if it isn't going to be used. If the first three `stdio` fields are marked as `UV_IGNORE` they'll redirect to `/dev/null`.

Since we want to pass on an existing descriptor, we'll use `UV_INHERIT_FD`. Then we set the `fd` to `stderr`.

## **proc-streams/main.c**

```
1
2 int main() {
3     loop = uv_default_loop();
4
5     /* ... */
6
7     options.stdio_count = 3;
8     uv_stdio_container_t child_stdio[3];
9     child_stdio[0].flags = UV_IGNORE;
10    child_stdio[1].flags = UV_IGNORE;
11    child_stdio[2].flags = UV_INHERIT_FD;
```

```

12     child_stdio[2].data.fd = 2;
13     options.stdio = child_stdio;
14
15     options.exit_cb = on_exit;
16     options.file = args[0];
17     options.args = args;
18
19     int r;
20     if ((r = uv_spawn(loop, &child_req, &options))) {
21         fprintf(stderr, "%s\n", uv_strerror(r));
22         return 1;
23     }
24
25     return uv_run(loop, UV_RUN_DEFAULT);
26 }

```

If you run `proc-stream` you'll see that only the line "This is stderr" will be displayed. Try marking `stdout` as being inherited and see the output.

It is dead simple to apply this redirection to streams. By setting `flags` to `UV_INHERIT_STREAM` and setting `data.stream` to the stream in the parent process, the child process can treat that stream as standard I/O. This can be used to implement something like [CGI](#)

[[https://en.wikipedia.org/wiki/Common\\_Gateway\\_Interface](https://en.wikipedia.org/wiki/Common_Gateway_Interface)].

A sample CGI script/executable is:

### cgi/tick.c

```

#include <stdio.h>
#include <unistd.h>

int main() {
    int i;
    for (i = 0; i < 10; i++) {
        printf("tick\n");
        fflush(stdout);
        sleep(1);
    }
}

```



```
}
printf("BOOM!\n");
return 0;
}
```

The CGI server combines the concepts from this chapter and [Networking](#) so that every client is sent ten ticks after which that connection is closed.

### cgi/main.c

```
1 void on_new_connection(uv_stream_t *server, int status)
2 {
3     if (status == -1) {
4         // error!
5         return;
6     }
7
8     uv_tcp_t *client = (uv_tcp_t*)
9     malloc(sizeof(uv_tcp_t));
10    uv_tcp_init(loop, client);
11    if (uv_accept(server, (uv_stream_t*) client) == 0)
12    {
13        invoke_cgi_script(client);
14    }
15    else {
16        uv_close((uv_handle_t*) client, NULL);
17    }
18 }
```

Here we simply accept the TCP connection and pass on the socket (*stream*) to `invoke_cgi_script`.

### cgi/main.c

```
1
2     args[1] = NULL;
3
4     /* ... finding the executable path and setting up
```

```

5  arguments ... */
6
7      options.stdio_count = 3;
8      uv_stdio_container_t child_stdio[3];
9      child_stdio[0].flags = UV_IGNORE;
10     child_stdio[1].flags = UV_INHERIT_STREAM;
11     child_stdio[1].data.stream = (uv_stream_t*) client;
12     child_stdio[2].flags = UV_IGNORE;
13     options.stdio = child_stdio;
14
15     options.exit_cb = cleanup_handles;
16     options.file = args[0];
17     options.args = args;
18
19     // Set this so we can close the socket after the
20     child process exits.
21     child_req.data = (void*) client;
22     int r;
23     if ((r = uv_spawn(loop, &child_req, &options))) {
24         fprintf(stderr, "%s\n", uv_strerror(r));

```

The `stdout` of the CGI script is set to the socket so that whatever our tick script prints, gets sent to the client. By using processes, we can offload the read/write buffering to the operating system, so in terms of convenience this is great. Just be warned that creating processes is a costly task.

## Parent-child IPC

A parent and child can have one or two way communication over a pipe created by settings `uv_stdio_container_t.flags` to a bit-wise combination of `UV_CREATE_PIPE` and `UV_READABLE_PIPE` or `UV_WRITABLE_PIPE`. The read/write flag is from the perspective of the child process. In this case, the `uv_stream_t*` `stream` field must be set to point to an initialized, unopened `uv_pipe_t` instance.

## New stdio Pipes

The `uv_pipe_t` structure represents more than just [pipe\(7\)](https://man7.org/linux/man-pages/man7/pipe.7.html) [https://man7.org/linux/man-pages/man7/pipe.7.html] (or `|`), but supports any streaming file-like objects. On Windows, the only object of that description is the [Named Pipe](https://docs.microsoft.com/en-us/windows/win32/ipc/named-pipes) [https://docs.microsoft.com/en-us/windows/win32/ipc/named-pipes]. On Unix, this could be any of [Unix Domain Socket](https://man7.org/linux/man-pages/man7/unix.7.html) [https://man7.org/linux/man-pages/man7/unix.7.html], or derived from [mkfifo\(1\)](https://man7.org/linux/man-pages/man1/mkfifo.1.html) [https://man7.org/linux/man-pages/man1/mkfifo.1.html], or it could actually be a [pipe\(7\)](https://man7.org/linux/man-pages/man7/pipe.7.html) [https://man7.org/linux/man-pages/man7/pipe.7.html]. When `uv_spawn` initializes a `uv_pipe_t` due to the `UV_CREATE_PIPE` flag, it opts for creating a [socketpair\(2\)](https://man7.org/linux/man-pages/man2/socketpair.2.html) [https://man7.org/linux/man-pages/man2/socketpair.2.html].

This is intended for the purpose of allowing multiple libuv processes to communicate with IPC. This is discussed below.

## Arbitrary process IPC

Since domain sockets [\[1\]](https://man7.org/linux/man-pages/man2/socket.2.html) can have a well known name and a location in the file-system they can be used for IPC between unrelated processes. The [D-BUS](https://www.freedesktop.org/wiki/Software/dbus) [https://www.freedesktop.org/wiki/Software/dbus] system used by open source desktop environments uses domain sockets for event notification. Various applications can then react when a contact comes online or new hardware is detected. The MySQL server also runs a domain socket on which clients can interact with it.

When using domain sockets, a client-server pattern is usually followed with the creator/owner of the socket acting as the server. After the initial setup, messaging is no different from TCP, so we'll re-use the echo server example.

### pipe-echo-server/main.c

---

```

1 void remove_sock(int sig) {
2     uv_fs_t req;
3     uv_fs_unlink(loop, &req, PIPENAME, NULL);
4     exit(0);
5 }
6
7 int main() {
8     loop = uv_default_loop();
9
10    uv_pipe_t server;
11    uv_pipe_init(loop, &server, 0);
12
13    signal(SIGINT, remove_sock);
14
15    int r;
16    if ((r = uv_pipe_bind(&server, PIPENAME))) {
17        fprintf(stderr, "Bind error %s\n",
18 uv_err_name(r));
19        return 1;
20    }
21    if ((r = uv_listen((uv_stream_t*) &server, 128,
22 on_new_connection))) {
23        fprintf(stderr, "Listen error %s\n",
24 uv_err_name(r));
25        return 2;
26    }
27    return uv_run(loop, UV_RUN_DEFAULT);
28 }

```

We name the socket `echo.sock` which means it will be created in the local directory. This socket now behaves no different from TCP sockets as far as the stream API is concerned. You can test this server using [socat](http://www.dest-unreach.org/socat/) [http://www.dest-unreach.org/socat/]:

```
$ socat - /path/to/socket
```

A client which wants to connect to a domain socket will use:

```
void uv_pipe_connect(uv_connect_t *req, uv_pipe_t *handle,
const char *name, uv_connect_cb cb);
```

where `name` will be `echo.sock` or similar. On Unix systems, `name` must point to a valid file (e.g. `/tmp/echo.sock`). On Windows, `name` follows a `\\?\pipe\echo.sock` format.

## Sending file descriptors over pipes

The cool thing about domain sockets is that file descriptors can be exchanged between processes by sending them over a domain socket. This allows processes to hand off their I/O to other processes. Applications include load-balancing servers, worker processes and other ways to make optimum use of CPU. libuv only supports sending **TCP sockets or other pipes** over pipes for now.

To demonstrate, we will look at a echo server implementation that hands off clients to worker processes in a round-robin fashion. This program is a bit involved, and while only snippets are included in the book, it is recommended to read the full code to really understand it.

The worker process is quite simple, since the file-descriptor is handed over to it by the master.

### multi-echo-server/worker.c

```
1 uv_loop_t *loop;
2 uv_pipe_t queue;
3 int main() {
4     loop = uv_default_loop();
5
6     uv_pipe_init(loop, &queue, 1 /* ipc */);
7     uv_pipe_open(&queue, 0);
8     uv_read_start((uv_stream_t*)&queue, alloc_buffer,
9 on_new_connection);
10     return uv_run(loop, UV_RUN_DEFAULT);
11 }
```

queue is the pipe connected to the master process on the other end, along which new file descriptors get sent. It is important to set the `ipc` argument of `uv_pipe_init` to 1 to indicate this pipe will be used for inter-process communication! Since the master will write the file handle to the standard input of the worker, we connect the pipe to `stdin` using `uv_pipe_open`.

## multi-echo-server/worker.c

```
1 void on_new_connection(uv_stream_t *q, ssize_t nread,
2 const uv_buf_t *buf) {
3     if (nread < 0) {
4         if (nread != UV_EOF)
5             fprintf(stderr, "Read error %s\n",
6 uv_err_name(nread));
7         uv_close((uv_handle_t*) q, NULL);
8         return;
9     }
10
11     uv_pipe_t *pipe = (uv_pipe_t*) q;
12     if (!uv_pipe_pending_count(pipe)) {
13         fprintf(stderr, "No pending count\n");
14         return;
15     }
16
17     uv_handle_type pending =
18 uv_pipe_pending_type(pipe);
19     assert(pending == UV_TCP);
20
21     uv_tcp_t *client = (uv_tcp_t*)
22 malloc(sizeof(uv_tcp_t));
23     uv_tcp_init(loop, client);
24     if (uv_accept(q, (uv_stream_t*) client) == 0) {
25         uv_os_fd_t fd;
26         uv_fileno((const uv_handle_t*) client, &fd);
27         fprintf(stderr, "Worker %d: Accepted fd %d\n",
28 getpid(), fd);
29         uv_read_start((uv_stream_t*) client,
30 alloc_buffer, echo_read);
31     }
32 }
```

```
    else {
        uv_close((uv_handle_t*) client, NULL);
    }
}
```

First we call `uv_pipe_pending_count()` to ensure that a handle is available to read out. If your program could deal with different types of handles, `uv_pipe_pending_type()` can be used to determine the type. Although `accept` seems odd in this code, it actually makes sense. What `accept` traditionally does is get a file descriptor (the `client`) from another file descriptor (The listening socket). Which is exactly what we do here. Fetch the file descriptor (`client`) from `queue`. From this point the worker does standard echo server stuff.

Turning now to the master, let's take a look at how the workers are launched to allow load balancing.

### **multi-echo-server/main.c**

```
1  struct child_worker {
2      uv_process_t req;
3      uv_process_options_t options;
4      uv_pipe_t pipe;
5  } *workers;
```

The `child_worker` structure wraps the process, and the pipe between the master and the individual process.

### **multi-echo-server/main.c**

```
1  void setup_workers() {
2      round_robin_counter = 0;
3
4      // ...
```

```

5
6 // launch same number of workers as number of CPUs
7 uv_cpu_info_t *info;
8 int cpu_count;
9 uv_cpu_info(&info, &cpu_count);
10 uv_free_cpu_info(info, cpu_count);
11
12 child_worker_count = cpu_count;
13
14 workers = calloc(cpu_count, sizeof(struct
15 child_worker));
16 while (cpu_count--) {
17     struct child_worker *worker =
18 &workers[cpu_count];
19     uv_pipe_init(loop, &worker->pipe, 1);
20
21     uv_stdio_container_t child_stdio[3];
22     child_stdio[0].flags = UV_CREATE_PIPE |
23 UV_READABLE_PIPE;
24     child_stdio[0].data.stream = (uv_stream_t*)
25 &worker->pipe;
26     child_stdio[1].flags = UV_IGNORE;
27     child_stdio[2].flags = UV_INHERIT_FD;
28     child_stdio[2].data.fd = 2;
29
30     worker->options.stdio = child_stdio;
31     worker->options.stdio_count = 3;
32
33     worker->options.exit_cb = close_process_handle;
34     worker->options.file = args[0];
35     worker->options.args = args;
36
37     uv_spawn(loop, &worker->req, &worker->options);
38     fprintf(stderr, "Started worker %d\n", worker-
39 >req.pid);
40 }
41 }

```

In setting up the workers, we use the nifty libuv function `uv_cpu_info` to get the number of CPUs so we can launch an equal number of workers. Again it is important to initialize the pipe acting as the IPC channel with the third argument as 1. We then indicate that the child process' `stdin` is to be a



readable pipe (from the point of view of the child). Everything is straightforward till here. The workers are launched and waiting for file descriptors to be written to their standard input.

It is in `on_new_connection` (the TCP infrastructure is initialized in `main()`), that we accept the client socket and pass it along to the next worker in the round-robin.

### multi-echo-server/main.c

```
void on_new_connection(uv_stream_t *server, int status)
{
    if (status == -1) {
        // error!
        return;
    }

    uv_tcp_t *client = (uv_tcp_t*)
    malloc(sizeof(uv_tcp_t));
    uv_tcp_init(loop, client);
    if (uv_accept(server, (uv_stream_t*) client) == 0)
    {
        uv_write_t *write_req = (uv_write_t*)
        malloc(sizeof(uv_write_t));
        dummy_buf = uv_buf_init("a", 1);
        struct child_worker *worker =
        &workers[round_robin_counter];
        uv_write2(write_req, (uv_stream_t*) &worker-
        >pipe, &dummy_buf, 1, (uv_stream_t*) client, NULL);
        round_robin_counter = (round_robin_counter + 1)
        % child_worker_count;
    }
    else {
        uv_close((uv_handle_t*) client, NULL);
    }
}
```

The `uv_write2` call handles all the abstraction and it is simply a matter of passing in the handle (`client`) as the right

argument. With this our multi-process echo server is operational.

Thanks to Kyle for [pointing out](#)

[<https://github.com/nikhilm/uvbook/issues/56>] that `uv_write2()` requires a non-empty buffer even when sending handles.

---

[[1](#)] In this section domain sockets stands in for named pipes on Windows as well.

# Advanced event loops

libuv provides considerable user control over event loops, and you can achieve interesting results by juggling multiple loops. You can also embed libuv's event loop into another event loop based library - imagine a Qt based UI, and Qt's event loop driving a libuv backend which does intensive system level tasks.

## Stopping an event loop

`uv_stop()` can be used to stop an event loop. The earliest the loop will stop running is *on the next iteration*, possibly later. This means that events that are ready to be processed in this iteration of the loop will still be processed, so `uv_stop()` can't be used as a kill switch. When `uv_stop()` is called, the loop **won't** block for i/o on this iteration. The semantics of these things can be a bit difficult to understand, so let's look at `uv_run()` where all the control flow occurs.

### src/unix/core.c - uv\_run

```
1         break;
2
3     case UV_UDP:
4         uv__udp_finish_close((uv_udp_t*)handle);
5         break;
6
7     default:
8         assert(0);
9         break;
10    }
11
12    uv__handle_unref(handle);
13    QUEUE_REMOVE(&handle->handle_queue);
```

```

14
15     if (handle->close_cb) {
16         handle->close_cb(handle);
17     }
18 }
19
20
21 static void uv__run_closing_handles(uv_loop_t* loop) {

```

stop\_flag is set by uv\_stop(). Now all libuv callbacks are invoked within the event loop, which is why invoking uv\_stop() in them will still lead to this iteration of the loop occurring. First libuv updates timers, then runs pending timer, idle and prepare callbacks, and invokes any pending I/O callbacks. If you were to call uv\_stop() in any of them, stop\_flag would be set. This causes uv\_backend\_timeout() to return 0, which is why the loop does not block on I/O. If on the other hand, you called uv\_stop() in one of the check handlers, I/O has already finished and is not affected.

uv\_stop() is useful to shutdown a loop when a result has been computed or there is an error, without having to ensure that all handlers are stopped one by one.

Here is a simple example that stops the loop and demonstrates how the current iteration of the loop still takes places.

### uvstop/main.c

```

1  #include <stdio.h>
2  #include <uv.h>
3
4  int64_t counter = 0;
5
6  void idle_cb(uv_idle_t *handle) {
7      printf("Idle callback\n");
      counter++;

```

```
8
9     if (counter >= 5) {
10         uv_stop(uv_default_loop());
11         printf("uv_stop() called\n");
12     }
13 }
14
15 void prep_cb(uv_prepare_t *handle) {
16     printf("Prep callback\n");
17 }
18
19 int main() {
20     uv_idle_t idler;
21     uv_prepare_t prep;
22
23     uv_idle_init(uv_default_loop(), &idler);
24     uv_idle_start(&idler, idle_cb);
25
26     uv_prepare_init(uv_default_loop(), &prep);
27     uv_prepare_start(&prep, prep_cb);
28
29     uv_run(uv_default_loop(), UV_RUN_DEFAULT);
30
31     return 0;
32 }
33
```

# Utilities

This chapter catalogues tools and techniques which are useful for common tasks. The [libev man page](#) [http://pod.tst.eu/http://cvs.schmorp.de/libev/ev.pod#COMMON\_OR\_USEFUL\_IDS\_OR\_BOTH] already covers some patterns which can be adopted to libuv through simple API changes. It also covers parts of the libuv API that don't require entire chapters dedicated to them.

## Timers

Timers invoke the callback after a certain time has elapsed since the timer was started. libuv timers can also be set to invoke at regular intervals instead of just once.

Simple use is to init a watcher and start it with a `timeout`, and optional `repeat`. Timers can be stopped at any time.

```
uv_timer_t timer_req;  
  
uv_timer_init(loop, &timer_req);  
uv_timer_start(&timer_req, callback, 5000, 2000);
```

will start a repeating timer, which first starts 5 seconds (the `timeout`) after the execution of `uv_timer_start`, then repeats every 2 seconds (the `repeat`). Use:

```
uv_timer_stop(&timer_req);
```

to stop the timer. This can be used safely from within the callback as well.

The repeat interval can be modified at any time with:

```
uv_timer_set_repeat(uv_timer_t *timer, int64_t repeat);
```

which will take effect **when possible**. If this function is called from a timer callback, it means:

- If the timer was non-repeating, the timer has already been stopped. Use `uv_timer_start` again.
- If the timer is repeating, the next timeout has already been scheduled, so the old repeat interval will be used once more before the timer switches to the new interval.

The utility function:

```
int uv_timer_again(uv_timer_t *)
```

applies **only to repeating timers** and is equivalent to stopping the timer and then starting it with both initial timeout and repeat set to the old repeat value. If the timer hasn't been started it fails (error code `UV_EINVAL`) and returns -1.

An actual timer example is in the [reference count section](#).

## Event loop reference count

The event loop only runs as long as there are active handles. This system works by having every handle increase the reference count of the event loop when it is started and decreasing the reference count when stopped. It is also possible to manually change the reference count of handles using:

```
void uv_ref(uv_handle_t*);  
void uv_unref(uv_handle_t*);
```

These functions can be used to allow a loop to exit even when a watcher is active or to use custom objects to keep the loop alive.

The latter can be used with interval timers. You might have a garbage collector which runs every X seconds, or your network service might send a heartbeat to others periodically, but you don't want to have to stop them along all clean exit paths or error scenarios. Or you want the program to exit when all your other watchers are done. In that case just unref the timer immediately after creation so that if it is the only watcher running then `uv_run` will still exit.

This is also used in `node.js` where some `libuv` methods are being bubbled up to the JS API. A `uv_handle_t` (the superclass of all watchers) is created per JS object and can be `ref/unref`ed.

## ref-timer/main.c

```
1 uv_loop_t *loop;
2 uv_timer_t gc_req;
3 uv_timer_t fake_job_req;
4
5 int main() {
6     loop = uv_default_loop();
7
8     uv_timer_init(loop, &gc_req);
9     uv_unref((uv_handle_t*) &gc_req);
10
11     uv_timer_start(&gc_req, gc, 0, 2000);
12
13     // could actually be a TCP download or something
14     uv_timer_init(loop, &fake_job_req);
15     uv_timer_start(&fake_job_req, fake_job, 9000, 0);
16     return uv_run(loop, UV_RUN_DEFAULT);
17 }
```



We initialize the garbage collector timer, then immediately unref it. Observe how after 9 seconds, when the fake job is done, the program automatically exits, even though the garbage collector is still running.

## Idler pattern

The callbacks of idle handles are invoked once per event loop. The idle callback can be used to perform some very low priority activity. For example, you could dispatch a summary of the daily application performance to the developers for analysis during periods of idleness, or use the application's CPU time to perform SETI calculations :) An idle watcher is also useful in a GUI application. Say you are using an event loop for a file download. If the TCP socket is still being established and no other events are present your event loop will pause (**block**), which means your progress bar will freeze and the user will face an unresponsive application. In such a case queue up an idle watcher to keep the UI operational.

### idle-compute/main.c

```
1 uv_loop_t *loop;
2 uv_fs_t stdin_watcher;
3 uv_idle_t idler;
4 char buffer[1024];
5
6 int main() {
7     loop = uv_default_loop();
8
9     uv_idle_init(loop, &idler);
10
11     uv_buf_t buf = uv_buf_init(buffer, 1024);
12     uv_fs_read(loop, &stdin_watcher, 0, &buf, 1, -1,
13 on_type);
14     uv_idle_start(&idler, crunch_away);
```

```
14     return uv_run(loop, UV_RUN_DEFAULT);
15 }
```

Here we initialize the idle watcher and queue it up along with the actual events we are interested in. `crunch_away` will now be called repeatedly until the user types something and presses Return. Then it will be interrupted for a brief amount as the loop deals with the input data, after which it will keep calling the idle callback again.

### idle-compute/main.c

```
1 void crunch_away(uv_idle_t* handle) {
2     // Compute extra-terrestrial life
3     // fold proteins
4     // computer another digit of PI
5     // or similar
6     fprintf(stderr, "Computing PI...\n");
7     // just to avoid overwhelming your terminal
8     emulator
9     uv_idle_stop(handle);
10 }
```

## Passing data to worker thread

When using `uv_queue_work` you'll usually need to pass complex data through to the worker thread. The solution is to use a struct and set `uv_work_t.data` to point to it. A slight variation is to have the `uv_work_t` itself as the first member of this struct (called a baton [\[1\]](#)). This allows cleaning up the work request and all the data in one free call.

```
1 struct ftp_baton {
2     uv_work_t req;
3     char *host;
```

```

3     int port;
4     char *username;
5     char *password;
6 }
7

```

```

ftp_baton *baton = (ftp_baton*)
1 malloc(sizeof(ftp_baton));
2 baton->req.data = (void*) baton;
3 baton->host = strdup("my.webhost.com");
4 baton->port = 21;
5 // ...
6
7 uv_queue_work(loop, &baton->req, ftp_session,
ftp_cleanup);

```

Here we create the baton and queue the task.

Now the task function can extract the data it needs:

```

1 void ftp_session(uv_work_t *req) {
2     ftp_baton *baton = (ftp_baton*) req->data;
3
4     fprintf(stderr, "Connecting to %s\n", baton->host);
5 }
6
7 void ftp_cleanup(uv_work_t *req) {
8     ftp_baton *baton = (ftp_baton*) req->data;
9
10    free(baton->host);
11    // ...
12    free(baton);
13 }

```

We then free the baton which also frees the watcher.

## External I/O with polling

Usually third-party libraries will handle their own I/O, and keep track of their sockets and other files internally. In this case it isn't possible to use the standard stream I/O operations, but the library can still be integrated into the libuv event loop. All that is required is that the library allow you to access the underlying file descriptors and provide functions that process tasks in small increments as decided by your application. Some libraries though will not allow such access, providing only a standard blocking function which will perform the entire I/O transaction and only then return. It is unwise to use these in the event loop thread, use the [Thread pool work scheduling](#) instead. Of course, this will also mean losing granular control on the library.

The `uv_poll` section of libuv simply watches file descriptors using the operating system notification mechanism. In some sense, all the I/O operations that libuv implements itself are also backed by `uv_poll` like code. Whenever the OS notices a change of state in file descriptors being polled, libuv will invoke the associated callback.

Here we will walk through a simple download manager that will use [libcurl](https://curl.haxx.se/libcurl/) to download files. Rather than give all control to libcurl, we'll instead be using the libuv event loop, and use the non-blocking, async [multi](https://curl.haxx.se/libcurl/c/libcurl-multi.html) interface to progress with the download whenever libuv notifies of I/O readiness.

## uvwget/main.c - The setup

```
1 #include <assert.h>
2 #include <stdio.h>
3 #include <stdlib.h>
4 #include <uv.h>
5 #include <curl/curl.h>
6
```

```

7 uv_loop_t *loop;
8 CURLM *curl_handle;
9 uv_timer_t timeout;
10 int main(int argc, char **argv) {
11     loop = uv_default_loop();
12
13     if (argc <= 1)
14         return 0;
15
16     if (curl_global_init(CURL_GLOBAL_ALL)) {
17         fprintf(stderr, "Could not init cURL\n");
18         return 1;
19     }
20
21     uv_timer_init(loop, &timeout);
22
23     curl_handle = curl_multi_init();
24     curl_multi_setopt(curl_handle,
25     CURLOPT_SOCKETFUNCTION, handle_socket);
26     curl_multi_setopt(curl_handle,
27     CURLOPT_TIMERFUNCTION, start_timeout);
28
29     while (argc-- > 1) {
30         add_download(argv[argc], argc);
31     }
32
33     uv_run(loop, UV_RUN_DEFAULT);
34     curl_multi_cleanup(curl_handle);
35     return 0;
36 }

```

The way each library is integrated with libuv will vary. In the case of libcurl, we can register two callbacks. The socket callback `handle_socket` is invoked whenever the state of a socket changes and we have to start polling it. `start_timeout` is called by libcurl to notify us of the next timeout interval, after which we should drive libcurl forward regardless of I/O status. This is so that libcurl can handle errors or do whatever else is required to get the download moving.

Our downloader is to be invoked as:

```
$ ./uvwget [url1] [url2] ...
```

So we add each argument as an URL

## uvwget/main.c - Adding urls

```
1 void add_download(const char *url, int num) {
2     char filename[50];
3     sprintf(filename, "%d.download", num);
4     FILE *file;
5
6     file = fopen(filename, "w");
7     if (file == NULL) {
8         fprintf(stderr, "Error opening %s\n",
9 filename);
10        return;
11    }
12
13    CURL *handle = curl_easy_init();
14    curl_easy_setopt(handle, CURLOPT_WRITEDATA, file);
15    curl_easy_setopt(handle, CURLOPT_URL, url);
16    curl_multi_add_handle(curl_handle, handle);
17    fprintf(stderr, "Added download %s -> %s\n", url,
18 filename);
19 }
```

We let libcurl directly write the data to a file, but much more is possible if you so desire.

`start_timeout` will be called immediately the first time by libcurl, so things are set in motion. This simply starts a [libuv timer](#) which drives `curl_multi_socket_action` with `CURL_SOCKET_TIMEOUT` whenever it times out.

`curl_multi_socket_action` is what drives libcurl, and what we call whenever sockets change state. But before we go into that, we need to poll on sockets whenever `handle_socket` is called.

## uvwget/main.c - Setting up polling

```
1 void start_timeout(CURLM *multi, long timeout_ms, void
2 *userp) {
3     if (timeout_ms <= 0)
4         timeout_ms = 1; /* 0 means directly call
5 socket_action, but we'll do it in a bit */
6     uv_timer_start(&timeout, on_timeout, timeout_ms,
7 0);
8 }
9
10
11 int handle_socket(CURL *easy, curl_socket_t s, int
12 action, void *userp, void *socketp) {
13     curl_context_t *curl_context;
14     if (action == CURL_POLL_IN || action ==
15 CURL_POLL_OUT) {
16         if (socketp) {
17             curl_context = (curl_context_t*) socketp;
18         }
19         else {
20             curl_context = create_curl_context(s);
21             curl_multi_assign(curl_handle, s, (void *)
22 curl_context);
23         }
24     }
25
26     switch (action) {
27     case CURL_POLL_IN:
28         uv_poll_start(&curl_context->poll_handle,
29 UV_READABLE, curl_perform);
30         break;
31     case CURL_POLL_OUT:
32         uv_poll_start(&curl_context->poll_handle,
33 UV_WRITABLE, curl_perform);
34         break;
35     case CURL_POLL_REMOVE:
36         if (socketp) {
37             uv_poll_stop(&
38 ((curl_context_t*)socketp)->poll_handle);
39             destroy_curl_context((curl_context_t*)
socketp);
40             curl_multi_assign(curl_handle, s,
NULL);

```

```

        }
        break;
    default:
        abort();
    }

    return 0;
}

```

We are interested in the socket `fd s`, and the `action`. For every socket we create a `uv_poll_t` handle if it doesn't exist, and associate it with the socket using `curl_multi_assign`. This way `socketp` points to it whenever the callback is invoked.

In the case that the download is done or fails, libcurl requests removal of the poll. So we stop and free the poll handle.

Depending on what events libcurl wishes to watch for, we start polling with `UV_READABLE` or `UV_WRITABLE`. Now libuv will invoke the poll callback whenever the socket is ready for reading or writing. Calling `uv_poll_start` multiple times on the same handle is acceptable, it will just update the events mask with the new value. `curl_perform` is the crux of this program.

### **uvwget/main.c - Driving libcurl.**

```

1 void curl_perform(uv_poll_t *req, int status, int
2 events) {
3     uv_timer_stop(&timeout);
4     int running_handles;
5     int flags = 0;
6     if (status < 0)                flags =
7     CURL_CSELECT_ERR;
8     if (!status && events & UV_READABLE) flags |=
9     CURL_CSELECT_IN;
10    if (!status && events & UV_WRITABLE) flags |=

```



```

11  CURL_CSELECT_OUT;
12
13     curl_context_t *context;
14
15     context = (curl_context_t*)req;

    curl_multi_socket_action(curl_handle, context-
>sockfd, flags, &running_handles);
    check_multi_info();
}

```

The first thing we do is to stop the timer, since there has been some progress in the interval. Then depending on what event triggered the callback, we set the correct flags. Then we call `curl_multi_socket_action` with the socket that progressed and the flags informing about what events happened. At this point libcurl does all of its internal tasks in small increments, and will attempt to return as fast as possible, which is exactly what an evented program wants in its main thread. libcurl keeps queueing messages into its own queue about transfer progress. In our case we are only interested in transfers that are completed. So we extract these messages, and clean up handles whose transfers are done.

### **uvwget/main.c - Reading transfer status.**

```

1  void check_multi_info(void) {
2      char *done_url;
3      CURLMsg *message;
4      int pending;
5
6      while ((message = curl_multi_info_read(curl_handle,
7 &pending))) {
8          switch (message->msg) {
9              case CURLMSG_DONE:
10                 curl_easy_getinfo(message->easy_handle,
11 CURLINFO_EFFECTIVE_URL,
                                     &done_url);

```

```

12         printf("%s DONE\n", done_url);
13
14         curl_multi_remove_handle(curl_handle,
15 message->easy_handle);
16         curl_easy_cleanup(message->easy_handle);
17         break;
18
19     default:
20         fprintf(stderr, "CURLMSG default\n");
21         abort();
22     }
}
}

```

## Check & Prepare watchers

TODO


## Loading libraries

libuv provides a cross platform API to dynamically load [shared libraries](#)

[[https://en.wikipedia.org/wiki/Shared\\_library#Shared\\_libraries](https://en.wikipedia.org/wiki/Shared_library#Shared_libraries)]. This can be used to implement your own plugin/extension/module system and is used by node.js to implement `require()` support for bindings. The usage is quite simple as long as your library exports the right symbols. Be careful with sanity and security checks when loading third party code, otherwise your program will behave unpredictably. This example implements a very simple plugin system which does nothing except print the name of the plugin.

Let us first look at the interface provided to plugin authors.

### plugin/plugin.h



```
1 #ifndef UVBOOK_PLUGIN_SYSTEM
2 #define UVBOOK_PLUGIN_SYSTEM
3
4 // Plugin authors should use this to register their
5 // plugins with mfp.
6 void mfp_register(const char *name);
7
8 #endif
```

You can similarly add more functions that plugin authors can use to do useful things in your application [\[2\]](#). A sample plugin using this API is:

### plugin/hello.c

```
1 #include "plugin.h"
2
3 void initialize() {
4     mfp_register("Hello World!");
5 }
```

Our interface defines that all plugins should have an `initialize` function which will be called by the application. This plugin is compiled as a shared library and can be loaded by running our application:

```
$ ./plugin libhello.dylib
Loading libhello.dylib
Registered plugin "Hello World!"
```

### Note

The shared library filename will be different depending on platforms. On Linux it is `libhello.so`.

This is done by using `uv_dlopen` to first load the shared library `libhello.dylib`. Then we get access to the `initialize` function using `uv_dlsym` and invoke it.

## plugin/main.c

```
1  #include "plugin.h"
2
3  typedef void (*init_plugin_function)();
4
5  void mfp_register(const char *name) {
6      fprintf(stderr, "Registered plugin \"%s\"\n",
7      name);
8  }
9
10 int main(int argc, char **argv) {
11     if (argc == 1) {
12         fprintf(stderr, "Usage: %s [plugin1] [plugin2]
13         ...\n", argv[0]);
14         return 0;
15     }
16
17     uv_lib_t *lib = (uv_lib_t*)
18     malloc(sizeof(uv_lib_t));
19     while (--argc) {
20         fprintf(stderr, "Loading %s\n", argv[argc]);
21         if (uv_dlopen(argv[argc], lib)) {
22             fprintf(stderr, "Error: %s\n",
23             uv_dlerror(lib));
24             continue;
25         }
26
27         init_plugin_function init_plugin;
28         if (uv_dlsym(lib, "initialize", (void **)
29         &init_plugin)) {
30             fprintf(stderr, "dlsym error: %s\n",
31             uv_dlerror(lib));
32             continue;
33         }
34
35         init_plugin();
36     }
37 }
```

```
    return 0;  
}
```

`uv_dlopen` expects a path to the shared library and sets the opaque `uv_lib_t` pointer. It returns 0 on success, -1 on error. Use `uv_dlerror` to get the error message.

`uv_dlsym` stores a pointer to the symbol in the second argument in the third argument. `init_plugin_function` is a function pointer to the sort of function we are looking for in the application's plugins.

## TTY

Text terminals have supported basic formatting for a long time, with a [pretty standardised](https://en.wikipedia.org/wiki/ANSI_escape_sequences)

[[https://en.wikipedia.org/wiki/ANSI\\_escape\\_sequences](https://en.wikipedia.org/wiki/ANSI_escape_sequences)] command set. This formatting is often used by programs to improve the readability of terminal output. For example `grep --colour`. libuv provides the `uv_tty_t` abstraction (a stream) and related functions to implement the ANSI escape codes across all platforms. By this I mean that libuv converts ANSI codes to the Windows equivalent, and provides functions to get terminal information.

The first thing to do is to initialize a `uv_tty_t` with the file descriptor it reads/writes from. This is achieved with:

```
int uv_tty_init(uv_loop_t*, uv_tty_t*, uv_file fd, int unused)
```

The `unused` parameter is now auto-detected and ignored. It previously needed to be set to use `uv_read_start()` on the stream.

It is then best to use `uv_tty_set_mode` to set the mode to *normal* which enables most TTY formatting, flow-control and other settings. [Other](#)

[[http://docs.libuv.org/en/v1.x/tty.html#c.uv\\_tty\\_mode\\_t](http://docs.libuv.org/en/v1.x/tty.html#c.uv_tty_mode_t)] modes are also available.

Remember to call `uv_tty_reset_mode` when your program exits to restore the state of the terminal. Just good manners. Another set of good manners is to be aware of redirection. If the user redirects the output of your command to a file, control sequences should not be written as they impede readability and `grep`. To check if the file descriptor is indeed a TTY, call `uv_guess_handle` with the file descriptor and compare the return value with `UV_TTY`.

Here is a simple example which prints white text on a red background:

### **tty/main.c**

```
1  #include <stdio.h>
2  #include <string.h>
3  #include <unistd.h>
4  #include <uv.h>
5
6  uv_loop_t *loop;
7  uv_tty_t tty;
8  int main() {
9      loop = uv_default_loop();
10
11     uv_tty_init(loop, &tty, STDOUT_FILENO, 0);
12     uv_tty_set_mode(&tty, UV_TTY_MODE_NORMAL);
13
14     if (uv_guess_handle(1) == UV_TTY) {
15         uv_write_t req;
16         uv_buf_t buf;
17         buf.base = "\033[41;37m";
18         buf.len = strlen(buf.base);
19         uv_write(&req, (uv_stream_t*) &tty, &buf, 1,
```

```

20 NULL);
21     }
22
23     uv_write_t req;
24     uv_buf_t buf;
25     buf.base = "Hello TTY\n";
26     buf.len = strlen(buf.base);
27     uv_write(&req, (uv_stream_t*) &tty, &buf, 1, NULL);
28     uv_tty_reset_mode();
29     return uv_run(loop, UV_RUN_DEFAULT);
}

```

The final TTY helper is `uv_tty_get_winsize()` which is used to get the width and height of the terminal and returns 0 on success. Here is a small program which does some animation using the function and character position escape codes.

### tty-gravity/main.c

```

1  #include <stdio.h>
2  #include <string.h>
3  #include <unistd.h>
4  #include <uv.h>
5
6  uv_loop_t *loop;
7  uv_tty_t tty;
8  uv_timer_t tick;
9  uv_write_t write_req;
10 int width, height;
11 int pos = 0;
12 char *message = " Hello TTY ";
13
14 void update(uv_timer_t *req) {
15     char data[500];
16
17     uv_buf_t buf;
18     buf.base = data;
19     buf.len = sprintf(data,
20 "\033[2J\033[H\033[%dB\033[%luC\033[42;37m%s",
21                               pos,

```

```

22         (unsigned long) (width-
23 strlen(message))/2,
24         message);
25     uv_write(&write_req, (uv_stream_t*) &tty, &buf, 1,
26 NULL);
27
28     pos++;
29     if (pos > height) {
30         uv_tty_reset_mode();
31         uv_timer_stop(&tick);
32     }
33 }
34
35 int main() {
36     loop = uv_default_loop();
37
38     uv_tty_init(loop, &tty, STDOUT_FILENO, 0);
39     uv_tty_set_mode(&tty, 0);
40
41     if (uv_tty_get_winsize(&tty, &width, &height)) {
42         fprintf(stderr, "Could not get TTY
43 information\n");
44         uv_tty_reset_mode();
45         return 1;
46     }
47
48     fprintf(stderr, "Width %d, height %d\n", width,
height);
    uv_timer_init(loop, &tick);
    uv_timer_start(&tick, update, 200, 200);
    return uv_run(loop, UV_RUN_DEFAULT);
}

```

The escape codes are:

## Code Meaning

---

2 J      Clear part of the screen, 2 is entire screen

---



## Code Meaning

---

H Moves cursor to certain position, default top-left

---

*n* B Moves cursor down by *n* lines

---

*n* C Moves cursor right by *n* columns

---

m Obeys string of display settings, in this case green background (40+2), white text (30+7)

---

As you can see this is very useful to produce nicely formatted output, or even console based arcade games if that tickles your fancy. For fancier control you can try [ncurses](https://www.gnu.org/software/ncurses/ncurses.html) [https://www.gnu.org/software/ncurses/ncurses.html].

*Changed in version 1.23.1::* the *readable* parameter is now unused and ignored. The appropriate value will now be auto-detected from the kernel.

---

[1] I was first introduced to the term baton in this context, in Konstantin Käfer's excellent slides on writing node.js bindings - <https://kkaefer.com/node-cpp-modules/#baton>

[2] mfp is My Fancy Plugin

# About

[Nikhil Marathe](https://nikhilism.com) [https://nikhilism.com] started writing this book one afternoon (June 16, 2012) when he didn't feel like programming. He had recently been stung by the lack of good documentation on libuv while working on [node-taglib](https://github.com/nikhilm/node-taglib) [https://github.com/nikhilm/node-taglib]. Although reference documentation was present, there were no comprehensive tutorials. This book is the output of that need and tries to be accurate. That said, the book may have mistakes. Pull requests are encouraged.

Nikhil is indebted to Marc Lehmann's comprehensive [man page](http://pod.tst.eu/http://cvs.schmorp.de/libev/ev.pod) [http://pod.tst.eu/http://cvs.schmorp.de/libev/ev.pod] about libev which describes much of the semantics of the two libraries.

This book was made using [Sphinx](https://www.sphinx-doc.org) [https://www.sphinx-doc.org] and [vim](https://www.vim.org) [https://www.vim.org].

## Note

In 2017 the libuv project incorporated the Nikhil's work into the official documentation and it's maintained there henceforth.

# Upgrading

Migration guides for different libuv versions, starting with 1.0.

- [libuv 0.10 -> 1.0.0 migration guide](#)

# libuv 0.10 -> 1.0.0 migration guide

Some APIs changed quite a bit throughout the 1.0.0 development process. Here is a migration guide for the most significant changes that happened after 0.10 was released.

## Loop initialization and closing

In libuv 0.10 (and previous versions), loops were created with *uv\_loop\_new*, which allocated memory for a new loop and initialized it; and destroyed with *uv\_loop\_delete*, which destroyed the loop and freed the memory. Starting with 1.0, those are deprecated and the user is responsible for allocating the memory and then initializing the loop.

libuv 0.10

```
uv_loop_t* loop = uv_loop_new();  
...  
uv_loop_delete(loop);
```

libuv 1.0

```
uv_loop_t* loop = malloc(sizeof *loop);  
uv_loop_init(loop);  
...  
uv_loop_close(loop);  
free(loop);
```

### Note

Error handling was omitted for brevity. Check the documentation for [uv\\_loop\\_init\(\)](#) and [uv\\_loop\\_close\(\)](#).

## Error handling

Error handling had a major overhaul in libuv 1.0. In general, functions and status parameters would get 0 for success and -1 for failure on libuv 0.10, and the user had to use `uv_last_error` to fetch the error code, which was a positive number.

In 1.0, functions and status parameters contain the actual error code, which is 0 for success, or a negative number in case of error.

libuv 0.10

```
... assume 'server' is a TCP server which is already listening
r = uv_listen((uv_stream_t*) server, 511, NULL);
if (r == -1) {
    uv_err_t err = uv_last_error(uv_default_loop());
    /* err.code contains UV_EADDRINUSE */
}
```

libuv 1.0

```
... assume 'server' is a TCP server which is already listening
r = uv_listen((uv_stream_t*) server, 511, NULL);
if (r < 0) {
    /* r contains UV_EADDRINUSE */
}
```

## Threadpool changes

In libuv 0.10 Unix used a threadpool which defaulted to 4 threads, while Windows used the `QueueUserWorkItem` API,

which uses a Windows internal threadpool, which defaults to 512 threads per process.

In 1.0, we unified both implementations, so Windows now uses the same implementation Unix does. The threadpool size can be set by exporting the `UV_THREADPOOL_SIZE` environment variable. See [Thread pool work scheduling](#).

## Allocation callback API change

In libuv 0.10 the callback had to return a filled `uv_buf_t` by value:

```
uv_buf_t alloc_cb(uv_handle_t* handle, size_t size) {  
    return uv_buf_init(malloc(size), size);  
}
```

In libuv 1.0 a pointer to a buffer is passed to the callback, which the user needs to fill:

```
void alloc_cb(uv_handle_t* handle, size_t size, uv_buf_t* buf)  
{  
    buf->base = malloc(size);  
    buf->len = size;  
}
```

## Unification of IPv4 / IPv6 APIs

libuv 1.0 unified the IPv4 and IPv6 APIs. There is no longer a `uv_tcp_bind` and `uv_tcp_bind6` duality, there is only [uv\\_tcp\\_bind\(\)](#) now.

IPv4 functions took `struct sockaddr_in` structures by value, and IPv6 functions took `struct sockaddr_in6`. Now functions take a `struct sockaddr*` (note it's a pointer). It can be stack allocated.

libuv 0.10

```
struct sockaddr_in addr = uv_ip4_addr("0.0.0.0", 1234);
...
uv_tcp_bind(&server, addr)
```

libuv 1.0

```
struct sockaddr_in addr;
uv_ip4_addr("0.0.0.0", 1234, &addr)
...
uv_tcp_bind(&server, (const struct sockaddr*) &addr, 0);
```

The IPv4 and IPv6 struct creating functions ([uv\\_ip4\\_addr\(\)](#) and [uv\\_ip6\\_addr\(\)](#)) have also changed, make sure you check the documentation.

..note::

This change applies to all functions that made a distinction between IPv4 and IPv6 addresses.

## Streams / UDP data receive callback API change

The streams and UDP data receive callbacks now get a pointer to a [uv\\_buf\\_t](#) buffer, not a structure by value.

libuv 0.10

```
void on_read(uv_stream_t* handle,
             ssize_t nread,
             uv_buf_t buf) {
    ...
}

void recv_cb(uv_udp_t* handle,
             ssize_t nread,
             uv_buf_t buf,
             struct sockaddr* addr,
             unsigned flags) {
```

```
    ...  
}
```

libuv 1.0

```
void on_read(uv_stream_t* handle,  
            ssize_t nread,  
            const uv_buf_t* buf) {  
    ...  
}  
  
void recv_cb(uv_udp_t* handle,  
            ssize_t nread,  
            const uv_buf_t* buf,  
            const struct sockaddr* addr,  
            unsigned flags) {  
    ...  
}
```

## Receiving handles over pipes API change

In libuv 0.10 (and earlier versions) the *uv\_read2\_start* function was used to start reading data on a pipe, which could also result in the reception of handles over it. The callback for such function looked like this:

```
void on_read(uv_pipe_t* pipe,  
            ssize_t nread,  
            uv_buf_t buf,  
            uv_handle_type pending) {  
    ...  
}
```

In libuv 1.0, *uv\_read2\_start* was removed, and the user needs to check if there are pending handles using [uv\\_pipe\\_pending\\_count\(\)](#) and [uv\\_pipe\\_pending\\_type\(\)](#) while in the read callback:



```

void on_read(uv_stream_t* handle,
             ssize_t nread,
             const uv_buf_t* buf) {
    ...
    while (uv_pipe_pending_count((uv_pipe_t*) handle) != 0) {
        pending = uv_pipe_pending_type((uv_pipe_t*) handle);
        ...
    }
    ...
}

```

## Extracting the file descriptor out of a handle

While it wasn't supported by the API, users often accessed the libuv internals in order to get access to the file descriptor of a TCP handle, for example.

```
fd = handle->io_watcher.fd;
```

This is now properly exposed through the [uv\\_fileno\(\)](#) function.

## uv\_fs\_readdir rename and API change

`uv_fs_readdir` returned a list of strings in the `req->ptr` field upon completion in libuv 0.10. In 1.0, this function got renamed to [uv\\_fs\\_scandir\(\)](#), since it's actually implemented using `scandir(3)`.

In addition, instead of allocating a full list strings, the user is able to get one result at a time by using the [uv\\_fs\\_scandir\\_next\(\)](#) function. This function does not need to make a roundtrip to the threadpool, because libuv will keep the list of *dents* returned by `scandir(3)` around.

# Index

## u

## U

[uv\\_accept \(C function\)](#),  
[uv\\_after\\_work\\_cb \(C type\)](#),  
[uv\\_alloc\\_cb \(C type\)](#),  
[uv\\_any\\_handle \(C type\)](#),  
[uv\\_any\\_req\\_ \(C type\)](#),  
[uv\\_async\\_cb \(C type\)](#),  
[uv\\_async\\_init \(C function\)](#),  
[uv\\_async\\_send \(C function\)](#),  
[uv\\_async\\_t \(C type\)](#),  
[uv\\_available\\_parallelism \(C function\)](#),  
[uv\\_backend\\_fd \(C function\)](#),  
[uv\\_backend\\_timeout \(C function\)](#),  
[uv\\_barrier\\_destroy \(C function\)](#),  
[uv\\_barrier\\_init \(C function\)](#),  
[uv\\_barrier\\_t \(C type\)](#),  
[uv\\_barrier\\_wait \(C function\)](#),  
[uv\\_buf\\_init \(C function\)](#),  
[uv\\_buf\\_t \(C type\)](#),  
[uv\\_buf\\_t.uv\\_buf\\_t.base \(C member\)](#),  
[uv\\_buf\\_t.uv\\_buf\\_t.len \(C member\)](#),  
[uv\\_calloc\\_func \(C type\)](#),  
[uv\\_cancel \(C function\)](#),  
[uv\\_chdir \(C function\)](#),  
[uv\\_check\\_cb \(C type\)](#),  
[uv\\_check\\_init \(C function\)](#),  
[uv\\_check\\_start \(C function\)](#),  
[uv\\_check\\_stop \(C function\)](#),  
[UV\\_IF\\_NAMESIZE \(C macro\)](#),  
[uv\\_inet\\_ntop \(C function\)](#),  
[uv\\_inet\\_pton \(C function\)](#),  
[uv\\_interface\\_address\\_t \(C type\)](#),  
[uv\\_interface\\_addresses \(C function\)](#),  
[uv\\_ip4\\_addr \(C function\)](#),  
[uv\\_ip4\\_name \(C function\)](#),  
[uv\\_ip6\\_addr \(C function\)](#),  
[uv\\_ip6\\_name \(C function\)](#),  
[uv\\_ip\\_name \(C function\)](#),  
[uv\\_is\\_active \(C function\)](#),  
[uv\\_is\\_closing \(C function\)](#),  
[uv\\_is\\_readable \(C function\)](#),  
[uv\\_is\\_writable \(C function\)](#),  
[uv\\_key\\_create \(C function\)](#),  
[uv\\_key\\_delete \(C function\)](#),  
[uv\\_key\\_get \(C function\)](#),  
[uv\\_key\\_set \(C function\)](#),  
[uv\\_key\\_t \(C type\)](#),  
[uv\\_kill \(C function\)](#),  
[uv\\_lib\\_t \(C type\)](#),  
[uv\\_listen \(C function\)](#),  
[uv\\_loadavg\\_ \(C function\)](#),  
[uv\\_loop\\_alive \(C function\)](#),  
[uv\\_loop\\_close \(C function\)](#),  
[uv\\_loop\\_configure \(C function\)](#),  
[uv\\_loop\\_fork \(C function\)](#),  
[uv\\_loop\\_get\\_data \(C function\)](#),  
[uv\\_loop\\_init \(C function\)](#),  
[uv\\_loop\\_set\\_data \(C function\)](#),  
[uv\\_loop\\_size \(C function\)](#),

[uv\\_check\\_t \(C type\)](#).  
[uv\\_close \(C function\)](#).  
[uv\\_close\\_cb \(C type\)](#).  
[uv\\_cond\\_broadcast \(C function\)](#).  
[uv\\_cond\\_destroy \(C function\)](#).  
[uv\\_cond\\_init \(C function\)](#).  
[uv\\_cond\\_signal \(C function\)](#).  
[uv\\_cond\\_t \(C type\)](#).  
[uv\\_cond\\_timedwait \(C function\)](#).  
[uv\\_cond\\_wait \(C function\)](#).  
[uv\\_connect\\_cb \(C type\)](#).  
[uv\\_connect\\_t \(C type\)](#).  
[uv\\_connect\\_t.handle \(C member\)](#).  
[uv\\_connection\\_cb \(C type\)](#).  
[uv\\_cpu\\_info \(C function\)](#).  
[uv\\_cpu\\_info\\_t \(C type\)](#).  
[uv\\_cwd \(C function\)](#).  
[uv\\_default\\_loop \(C function\)](#).  
[uv\\_dir\\_t \(C type\)](#).  
[uv\\_disable\\_stdio\\_inheritance \(C function\)](#).  
[uv\\_dlclose \(C function\)](#).  
[uv\\_dLError \(C function\)](#).  
[uv\\_dlopen \(C function\)](#).  
[uv\\_dlsym \(C function\)](#).  
[UV\\_E2BIG \(C macro\)](#).  
[UV\\_EACCES \(C macro\)](#).  
[UV\\_EADDRINUSE \(C macro\)](#).  
[UV\\_EADDRNOTAVAIL \(C macro\)](#).  
[UV\\_EAFNOSUPPORT \(C macro\)](#).  
[UV\\_EAGAIN \(C macro\)](#).  
[UV\\_EAI\\_ADDRFAMILY \(C macro\)](#).  
[UV\\_EAI\\_AGAIN \(C macro\)](#).

[uv\\_loop\\_t \(C type\)](#).  
[uv\\_loop\\_t.data \(C member\)](#).  
[uv\\_malloc\\_func \(C type\)](#).  
[uv\\_metrics\\_idle\\_time \(C function\)](#).  
[uv\\_mutex\\_destroy \(C function\)](#).  
[uv\\_mutex\\_init \(C function\)](#).  
[uv\\_mutex\\_init\\_recursive \(C function\)](#).  
[uv\\_mutex\\_lock \(C function\)](#).  
[uv\\_mutex\\_t \(C type\)](#).  
[uv\\_mutex\\_trylock \(C function\)](#).  
[uv\\_mutex\\_unlock \(C function\)](#).  
[uv\\_now \(C function\)](#).  
[uv\\_once \(C function\)](#).  
[uv\\_once\\_t \(C type\)](#).  
[uv\\_open\\_oshandle \(C function\)](#).  
[uv\\_os\\_environ \(C function\)](#).  
[uv\\_os\\_fd\\_t \(C type\)](#).  
[uv\\_os\\_free\\_passwd \(C function\)](#).  
[uv\\_os\\_get\\_passwd \(C function\)](#).  
[uv\\_os\\_getenv \(C function\)](#).  
[uv\\_os\\_gethostname \(C function\)](#).  
[uv\\_os\\_getpid \(C function\)](#).  
[uv\\_os\\_getppid \(C function\)](#).  
[uv\\_os\\_getpriority \(C function\)](#).  
[uv\\_os\\_homedir \(C function\)](#).  
[uv\\_os\\_setenv \(C function\)](#).  
[uv\\_os\\_setpriority \(C function\)](#).  
[uv\\_os\\_sock\\_t \(C type\)](#).  
[uv\\_os\\_tmpdir \(C function\)](#).  
[uv\\_os\\_uname \(C function\)](#).  
[uv\\_os\\_unsetenv \(C function\)](#).  
[uv\\_passwd\\_t \(C type\)](#).  
[uv\\_pid\\_t \(C type\)](#).  
[uv\\_pipe \(C function\)](#).  
[uv\\_pipe\\_bind \(C function\)](#).  
[uv\\_pipe\\_chmod \(C function\)](#).  
[uv\\_pipe\\_connect \(C function\)](#).  
[uv\\_pipe\\_getpeername \(C function\)](#).

[UV\\_EAI\\_BADFLAGS \(C macro\)](#),  
[UV\\_EAI\\_BADHINTS \(C macro\)](#),  
[UV\\_EAI\\_CANCELED \(C macro\)](#),  
[UV\\_EAI\\_FAIL \(C macro\)](#),  
[UV\\_EAI\\_FAMILY \(C macro\)](#),  
[UV\\_EAI\\_MEMORY \(C macro\)](#),  
[UV\\_EAI\\_NODATA \(C macro\)](#),  
[UV\\_EAI\\_NONAME \(C macro\)](#),  
[UV\\_EAI\\_OVERFLOW \(C macro\)](#),  
[UV\\_EAI\\_PROTOCOL \(C macro\)](#),  
[UV\\_EAI\\_SERVICE \(C macro\)](#),  
[UV\\_EAI\\_SOCKETYPE \(C macro\)](#),  
[UV\\_EALREADY \(C macro\)](#),  
[UV\\_EBADF \(C macro\)](#),  
[UV\\_EBUSY \(C macro\)](#),  
[UV\\_ECANCELED \(C macro\)](#),  
[UV\\_ECHARSET \(C macro\)](#),  
[UV\\_ECONNABORTED \(C macro\)](#),  
[UV\\_ECONNREFUSED \(C macro\)](#),  
[UV\\_ECONNRESET \(C macro\)](#),  
[UV\\_EDESTADDRREQ \(C macro\)](#),  
[UV\\_EEXIST \(C macro\)](#),  
[UV\\_EFAULT \(C macro\)](#),  
[UV\\_EFBIG \(C macro\)](#),  
[UV\\_EFTYPE \(C macro\)](#),  
[UV\\_EHOSTUNREACH \(C macro\)](#),  
[UV\\_EILSEQ \(C macro\)](#),  
[UV\\_EINTR \(C macro\)](#),  
[UV\\_EINVAL \(C macro\)](#),  
[UV\\_EIO \(C macro\)](#),  
[UV\\_EISCONN \(C macro\)](#),  
[UV\\_EISDIR \(C macro\)](#),  
[uv\\_pipe\\_getsockname \(C function\)](#),  
[uv\\_pipe\\_init \(C function\)](#),  
[uv\\_pipe\\_open \(C function\)](#),  
[uv\\_pipe\\_pending\\_count \(C function\)](#),  
[uv\\_pipe\\_pending\\_instances \(C function\)](#),  
[uv\\_pipe\\_pending\\_type \(C function\)](#),  
[uv\\_pipe\\_t \(C type\)](#),  
[uv\\_pipe\\_t.ipc \(C member\)](#),  
[uv\\_poll\\_cb \(C type\)](#),  
[uv\\_poll\\_event \(C type\)](#),  
[uv\\_poll\\_init \(C function\)](#),  
[uv\\_poll\\_init\\_socket \(C function\)](#),  
[uv\\_poll\\_start \(C function\)](#),  
[uv\\_poll\\_stop \(C function\)](#),  
[uv\\_poll\\_t \(C type\)](#),  
[uv\\_prepare\\_cb \(C type\)](#),  
[uv\\_prepare\\_init \(C function\)](#),  
[uv\\_prepare\\_start \(C function\)](#),  
[uv\\_prepare\\_stop \(C function\)](#),  
[uv\\_prepare\\_t \(C type\)](#),  
[uv\\_print\\_active\\_handles \(C function\)](#),  
[uv\\_print\\_all\\_handles \(C function\)](#),  
[uv\\_process\\_flags \(C type\)](#),  
[uv\\_process\\_get\\_pid \(C function\)](#),  
[uv\\_process\\_kill \(C function\)](#),  
[uv\\_process\\_options\\_t \(C type\)](#),  
[uv\\_process\\_options\\_t.args \(C member\)](#),  
[uv\\_process\\_options\\_t.cwd \(C member\)](#),  
[uv\\_process\\_options\\_t.env \(C member\)](#),  
[uv\\_process\\_options\\_t.exit\\_cb \(C member\)](#),  
[uv\\_process\\_options\\_t.file \(C member\)](#)

[UV\\_ELOOP \(C macro\)](#),  
[UV\\_EMFILE \(C macro\)](#),  
[UV\\_EMLINK \(C macro\)](#),  
[UV\\_MSGSIZE \(C macro\)](#),  
[UV\\_ENAMETOOLONG \(C macro\)](#),  
[UV\\_ENETDOWN \(C macro\)](#),  
[UV\\_ENETUNREACH \(C macro\)](#),  
[UV\\_ENFILE \(C macro\)](#),  
[UV\\_ENOBUFS \(C macro\)](#),  
[UV\\_ENODEV \(C macro\)](#),  
[UV\\_ENOENT \(C macro\)](#),  
[UV\\_ENOMEM \(C macro\)](#),  
[UV\\_ENONET \(C macro\)](#),  
[UV\\_ENOPROTOPT \(C macro\)](#),  
[UV\\_ENOSPC \(C macro\)](#),  
[UV\\_ENOSYS \(C macro\)](#),  
[UV\\_ENOTCONN \(C macro\)](#),  
[UV\\_ENOTDIR \(C macro\)](#),  
[UV\\_ENOTEMPTY \(C macro\)](#),  
[UV\\_ENOTSOCK \(C macro\)](#),  
[UV\\_ENOTSUP \(C macro\)](#),  
[UV\\_ENOTTY \(C macro\)](#),  
[uv\\_env\\_item\\_t \(C type\)](#),  
[UV\\_ENXIO \(C macro\)](#),  
[UV\\_EOF \(C macro\)](#),  
[UV\\_EOVERFLOW \(C macro\)](#),  
[UV\\_EPERM \(C macro\)](#),  
[UV\\_EPIPE \(C macro\)](#),  
[UV\\_EPROTO \(C macro\)](#),  
[UV\\_EPROTONOSUPPORT \(C macro\)](#),  
[UV\\_EPROTOTYPE \(C macro\)](#),  
[UV\\_ERANGE \(C macro\)](#),  
[UV\\_EROFS \(C macro\)](#),  
[uv\\_err\\_name \(C function\)](#),  
[uv\\_err\\_name\\_r \(C function\)](#),  
[UV\\_ERRNO\\_MAP \(C macro\)](#),  
[UV\\_ESHUTDOWN \(C macro\)](#),  
[uv\\_process\\_options\\_t.flags \(C member\)](#),  
[uv\\_process\\_options\\_t.gid \(C member\)](#),  
[uv\\_process\\_options\\_t.stdio \(C member\)](#),  
[uv\\_process\\_options\\_t.stdio\\_count \(C member\)](#),  
[uv\\_process\\_options\\_t.uid \(C member\)](#),  
[uv\\_process\\_t \(C type\)](#),  
[uv\\_process\\_t.pid \(C member\)](#),  
[uv\\_queue\\_work \(C function\)](#),  
[uv\\_random \(C function\)](#),  
[uv\\_random\\_cb \(C type\)](#),  
[uv\\_random\\_t \(C type\)](#),  
[uv\\_read\\_cb \(C type\)](#),  
[uv\\_read\\_start \(C function\)](#),  
[uv\\_read\\_stop \(C function\)](#),  
[uv\\_realloc\\_func \(C type\)](#),  
[uv\\_recv\\_buffer\\_size \(C function\)](#),  
[uv\\_ref \(C function\)](#),  
[uv\\_replace\\_allocator \(C function\)](#),  
[uv\\_req\\_get\\_data \(C function\)](#),  
[uv\\_req\\_get\\_type \(C function\)](#),  
[uv\\_req\\_set\\_data \(C function\)](#),  
[uv\\_req\\_size \(C function\)](#),  
[uv\\_req\\_t \(C type\)](#),  
[uv\\_req\\_t.data \(C member\)](#),  
[uv\\_req\\_t.type \(C member\)](#),  
[UV\\_REQ\\_TYPE\\_MAP \(C macro\)](#),  
[uv\\_req\\_type\\_name \(C function\)](#),  
[uv\\_resident\\_set\\_memory \(C function\)](#),  
[uv\\_run \(C function\)](#),  
[uv\\_rusage\\_t \(C type\)](#),  
[uv\\_rwlock\\_destroy \(C function\)](#),  
[uv\\_rwlock\\_init \(C function\)](#),  
[uv\\_rwlock\\_rdlock \(C function\)](#),  
[uv\\_rwlock\\_rdlunlock \(C function\)](#),  
[uv\\_rwlock\\_t \(C type\)](#)

[UV\\_ESOCKTNOSUPPORT \(C macro\)](#),  
[UV\\_ESPIPE \(C macro\)](#),  
[UV\\_ESRCH \(C macro\)](#),  
[UV\\_ETIMEDOUT \(C macro\)](#),  
[UV\\_ETXTBSY \(C macro\)](#),  
[UV\\_EXDEV \(C macro\)](#),  
[uv\\_exepath \(C function\)](#),  
[uv\\_exit\\_cb \(C type\)](#),  
[uv\\_file \(C type\)](#),  
[uv\\_fileno \(C function\)](#),  
[uv\\_free\\_cpu\\_info \(C function\)](#),  
[uv\\_free\\_func \(C type\)](#),  
[uv\\_free\\_interface\\_addresses \(C function\)](#),  
[uv\\_freeaddrinfo \(C function\)](#),  
[uv\\_fs\\_access \(C function\)](#),  
[uv\\_fs\\_chmod \(C function\)](#),  
[uv\\_fs\\_chown \(C function\)](#),  
[uv\\_fs\\_close \(C function\)](#),  
[uv\\_fs\\_closedir \(C function\)](#),  
[uv\\_fs\\_copyfile \(C function\)](#),  
[uv\\_fs\\_event \(C type\)](#),  
[uv\\_fs\\_event\\_cb \(C type\)](#),  
[uv\\_fs\\_event\\_flags \(C type\)](#),  
[uv\\_fs\\_event\\_getpath \(C function\)](#),  
[uv\\_fs\\_event\\_init \(C function\)](#),  
[uv\\_fs\\_event\\_start \(C function\)](#),  
[uv\\_fs\\_event\\_stop \(C function\)](#),  
[uv\\_fs\\_event\\_t \(C type\)](#),  
[uv\\_fs\\_fchmod \(C function\)](#),  
[uv\\_fs\\_fchown \(C function\)](#),  
[uv\\_fs\\_fdatasync \(C function\)](#),  
[uv\\_fs\\_fstat \(C function\)](#),  
[uv\\_fs\\_fsync \(C function\)](#),  
[uv\\_fs\\_ftruncate \(C function\)](#),  
[uv\\_fs\\_futime \(C function\)](#),  
[uv\\_rwlock\\_tryrdlock \(C function\)](#),  
[uv\\_rwlock\\_trywrlock \(C function\)](#),  
[uv\\_rwlock\\_wrlock \(C function\)](#),  
[uv\\_rwlock\\_wrunlock \(C function\)](#),  
[uv\\_sem\\_destroy \(C function\)](#),  
[uv\\_sem\\_init \(C function\)](#),  
[uv\\_sem\\_post \(C function\)](#),  
[uv\\_sem\\_t \(C type\)](#),  
[uv\\_sem\\_trywait \(C function\)](#),  
[uv\\_sem\\_wait \(C function\)](#),  
[uv\\_send\\_buffer\\_size \(C function\)](#),  
[uv\\_set\\_process\\_title \(C function\)](#),  
[uv\\_setup\\_args \(C function\)](#),  
[uv\\_shutdown \(C function\)](#),  
[uv\\_shutdown\\_cb \(C type\)](#),  
[uv\\_shutdown\\_t \(C type\)](#),  
[uv\\_shutdown\\_t.handle \(C member\)](#),  
[uv\\_signal\\_cb \(C type\)](#),  
[uv\\_signal\\_init \(C function\)](#),  
[uv\\_signal\\_start \(C function\)](#),  
[uv\\_signal\\_start\\_oneshot \(C function\)](#),  
[uv\\_signal\\_stop \(C function\)](#),  
[uv\\_signal\\_t \(C type\)](#),  
[uv\\_signal\\_t.signum \(C member\)](#),  
[uv\\_sleep \(C function\)](#),  
[uv\\_socketpair \(C function\)](#),  
[uv\\_spawn \(C function\)](#),  
[uv\\_stat\\_t \(C type\)](#),  
[uv\\_statfs\\_t \(C type\)](#),  
[uv\\_stdio\\_container\\_t \(C type\)](#),  
[uv\\_stdio\\_container\\_t.data \(C member\)](#),  
[uv\\_stdio\\_container\\_t.flags \(C member\)](#),  
[uv\\_stop \(C function\)](#),  
[uv\\_stream\\_get\\_write\\_queue\\_size \(C function\)](#),  
[uv\\_stream\\_set\\_blocking \(C function\)](#)

[uv\\_fs\\_get\\_path \(C function\)](#),  
[uv\\_fs\\_get\\_ptr \(C function\)](#),  
[uv\\_fs\\_get\\_result \(C function\)](#),  
[uv\\_fs\\_get\\_statbuf \(C function\)](#),  
[uv\\_fs\\_get\\_system\\_error \(C function\)](#),  
[uv\\_fs\\_get\\_type \(C function\)](#),  
[uv\\_fs\\_lchown \(C function\)](#),  
[uv\\_fs\\_link \(C function\)](#),  
[uv\\_fs\\_lstat \(C function\)](#),  
[uv\\_fs\\_lutime \(C function\)](#),  
[uv\\_fs\\_mkdir \(C function\)](#),  
[uv\\_fs\\_mkdtemp \(C function\)](#),  
[uv\\_fs\\_mkstemp \(C function\)](#),  
[UV\\_FS\\_O\\_APPEND \(C macro\)](#),  
[UV\\_FS\\_O\\_CREAT \(C macro\)](#),  
[UV\\_FS\\_O\\_DIRECT \(C macro\)](#),  
[UV\\_FS\\_O\\_DIRECTORY \(C macro\)](#),  
[UV\\_FS\\_O\\_DSYNC \(C macro\)](#),  
[UV\\_FS\\_O\\_EXCL \(C macro\)](#),  
[UV\\_FS\\_O\\_EXLOCK \(C macro\)](#),  
[UV\\_FS\\_O\\_FILEMAP \(C macro\)](#),  
[UV\\_FS\\_O\\_NOATIME \(C macro\)](#),  
[UV\\_FS\\_O\\_NOCTTY \(C macro\)](#),  
[UV\\_FS\\_O\\_NOFOLLOW \(C macro\)](#),  
[UV\\_FS\\_O\\_NONBLOCK \(C macro\)](#),  
[UV\\_FS\\_O\\_RANDOM \(C macro\)](#),  
[UV\\_FS\\_O\\_RDONLY \(C macro\)](#),  
[UV\\_FS\\_O\\_RDWR \(C macro\)](#),  
[UV\\_FS\\_O\\_SEQUENTIAL \(C macro\)](#),  
[UV\\_FS\\_O\\_SHORT\\_LIVED \(C macro\)](#),  
[UV\\_FS\\_O\\_SYMLINK \(C macro\)](#),

[uv\\_stream\\_t \(C type\)](#),  
[uv\\_stream\\_t.write\\_queue\\_size \(C member\)](#),  
[uv\\_strerror \(C function\)](#),  
[uv\\_strerror\\_r \(C function\)](#),  
[uv\\_tcp\\_bind \(C function\)](#),  
[uv\\_tcp\\_close\\_reset \(C function\)](#),  
[uv\\_tcp\\_connect \(C function\)](#),  
[uv\\_tcp\\_getpeername \(C function\)](#),  
[uv\\_tcp\\_getsockname \(C function\)](#),  
[uv\\_tcp\\_init \(C function\)](#),  
[uv\\_tcp\\_init\\_ex \(C function\)](#),  
[uv\\_tcp\\_keepalive \(C function\)](#),  
[uv\\_tcp\\_nodelay \(C function\)](#),  
[uv\\_tcp\\_open \(C function\)](#),  
[uv\\_tcp\\_simultaneous\\_accepts \(C function\)](#),  
[uv\\_tcp\\_t \(C type\)](#),  
[uv\\_thread\\_cb \(C type\)](#),  
[uv\\_thread\\_create \(C function\)](#),  
[uv\\_thread\\_create\\_ex \(C function\)](#),  
[uv\\_thread\\_equal \(C function\)](#),  
[uv\\_thread\\_join \(C function\)](#),  
[uv\\_thread\\_options\\_t \(C type\)](#),  
[uv\\_thread\\_self \(C function\)](#),  
[uv\\_thread\\_t \(C type\)](#),  
[uv\\_timer\\_again \(C function\)](#),  
[uv\\_timer\\_cb \(C type\)](#),  
[uv\\_timer\\_get\\_due\\_in \(C function\)](#),  
[uv\\_timer\\_get\\_repeat \(C function\)](#),  
[uv\\_timer\\_init \(C function\)](#),  
[uv\\_timer\\_set\\_repeat \(C function\)](#),  
[uv\\_timer\\_start \(C function\)](#),  
[uv\\_timer\\_stop \(C function\)](#),  
[uv\\_timer\\_t \(C type\)](#),  
[uv\\_timespec\\_t \(C type\)](#),  
[uv\\_timeval64\\_t \(C type\)](#),  
[uv\\_timeval\\_t \(C type\)](#),  
[uv\\_translate\\_sys\\_error \(C function\)](#),  
[uv\\_try\\_write \(C function\)](#),

[UV\\_FS\\_O\\_SYNC \(C macro\)](#)  
[UV\\_FS\\_O\\_TEMPORARY \(C macro\)](#)  
[UV\\_FS\\_O\\_TRUNC \(C macro\)](#)  
[UV\\_FS\\_O\\_WRONLY \(C macro\)](#)  
[uv\\_fs\\_open \(C function\)](#)  
[uv\\_fs\\_opendir \(C function\)](#)  
[uv\\_fs\\_poll\\_cb \(C type\)](#)  
[uv\\_fs\\_poll\\_getpath \(C function\)](#)  
[uv\\_fs\\_poll\\_init \(C function\)](#)  
[uv\\_fs\\_poll\\_start \(C function\)](#)  
[uv\\_fs\\_poll\\_stop \(C function\)](#)  
[uv\\_fs\\_poll\\_t \(C type\)](#)  
[uv\\_fs\\_read \(C function\)](#)  
[uv\\_fs\\_readdir \(C function\)](#)  
[uv\\_fs\\_readlink \(C function\)](#)  
[uv\\_fs\\_realpath \(C function\)](#)  
[uv\\_fs\\_rename \(C function\)](#)  
[uv\\_fs\\_req\\_cleanup \(C function\)](#)  
[uv\\_fs\\_rmdir \(C function\)](#)  
[uv\\_fs\\_scandir \(C function\)](#)  
[uv\\_fs\\_scandir\\_next \(C function\)](#)  
[uv\\_fs\\_sendfile \(C function\)](#)  
[uv\\_fs\\_stat \(C function\)](#)  
[uv\\_fs\\_statfs \(C function\)](#)  
[uv\\_fs\\_symlink \(C function\)](#)  
[uv\\_fs\\_t \(C type\)](#)  
[uv\\_fs\\_t.fs\\_type \(C member\)](#)  
[uv\\_fs\\_t.loop \(C member\)](#)  
[uv\\_fs\\_t.path \(C member\)](#)  
[uv\\_fs\\_t.ptr \(C member\)](#)  
[uv\\_fs\\_t.result \(C member\)](#)  
[uv\\_fs\\_t.statbuf \(C member\)](#)  
[uv\\_fs\\_unlink \(C function\)](#)  
[uv\\_fs\\_utime \(C function\)](#)  
[uv\\_fs\\_write \(C function\)](#)  
[uv\\_get\\_constrained\\_memory \(C function\)](#)

[uv\\_try\\_write2 \(C function\)](#)  
[uv\\_tty\\_get\\_vterm\\_state \(C function\)](#)  
[uv\\_tty\\_get\\_winsize \(C function\)](#)  
[uv\\_tty\\_init \(C function\)](#)  
[uv\\_tty\\_reset\\_mode \(C function\)](#)  
[uv\\_tty\\_set\\_mode \(C function\)](#)  
[uv\\_tty\\_set\\_vterm\\_state \(C function\)](#)  
[uv\\_tty\\_t \(C type\)](#)  
[uv\\_udp\\_bind \(C function\)](#)  
[uv\\_udp\\_connect \(C function\)](#)  
[uv\\_udp\\_flags \(C type\)](#)  
[uv\\_udp\\_get\\_send\\_queue\\_count \(C function\)](#)  
[uv\\_udp\\_get\\_send\\_queue\\_size \(C function\)](#)  
[uv\\_udp\\_getpeername \(C function\)](#)  
[uv\\_udp\\_getsockname \(C function\)](#)  
[uv\\_udp\\_init \(C function\)](#)  
[uv\\_udp\\_init\\_ex \(C function\)](#)  
[uv\\_udp\\_open \(C function\)](#)  
[uv\\_udp\\_recv\\_cb \(C type\)](#)  
[uv\\_udp\\_recv\\_start \(C function\)](#)  
[uv\\_udp\\_recv\\_stop \(C function\)](#)  
[uv\\_udp\\_send \(C function\)](#)  
[uv\\_udp\\_send\\_cb \(C type\)](#)  
[uv\\_udp\\_send\\_t \(C type\)](#)  
[uv\\_udp\\_send\\_t.handle \(C member\)](#)  
[uv\\_udp\\_set\\_broadcast \(C function\)](#)  
[uv\\_udp\\_set\\_membership \(C function\)](#)  
[uv\\_udp\\_set\\_multicast\\_interface \(C function\)](#)  
[uv\\_udp\\_set\\_multicast\\_loop \(C function\)](#)



[uv\\_get\\_free\\_memory](#) (C function),  
[uv\\_get\\_osfhandle](#) (C function),  
[uv\\_get\\_process\\_title](#) (C function),  
[uv\\_get\\_total\\_memory](#) (C function),  
[uv\\_getaddrinfo](#) (C function),  
[uv\\_getaddrinfo\\_cb](#) (C type),  
[uv\\_getaddrinfo\\_t](#) (C type),  
[uv\\_getaddrinfo\\_t.addrinfo](#) (C member),  
[uv\\_getaddrinfo\\_t.loop](#) (C member),  
[uv\\_getnameinfo](#) (C function),  
[uv\\_getnameinfo\\_cb](#) (C type),  
[uv\\_getnameinfo\\_t](#) (C type),  
[uv\\_getnameinfo\\_t.host](#) (C member),  
[uv\\_getnameinfo\\_t.loop](#) (C member),  
[uv\\_getnameinfo\\_t.service](#) (C member),  
[uv\\_getrusage](#) (C function),  
[uv\\_gettimeofday](#) (C function),  
[uv\\_guess\\_handle](#) (C function),  
[uv\\_handle\\_get\\_data](#) (C function),  
[uv\\_handle\\_get\\_loop](#) (C function),  
[uv\\_handle\\_get\\_type](#) (C function),  
[uv\\_handle\\_set\\_data](#) (C function),  
[uv\\_handle\\_size](#) (C function),  
[uv\\_handle\\_t](#) (C type),  
[uv\\_handle\\_t.data](#) (C member),  
[uv\\_udp\\_set\\_multicast\\_ttl](#) (C function),  
[uv\\_udp\\_set\\_source\\_membership](#) (C function),  
[uv\\_udp\\_set\\_ttl](#) (C function),  
[uv\\_udp\\_t](#) (C type),  
[uv\\_udp\\_t.send\\_queue\\_count](#) (C member),  
[uv\\_udp\\_t.send\\_queue\\_size](#) (C member),  
[uv\\_udp\\_try\\_send](#) (C function),  
[uv\\_udp\\_using\\_recvmmsg](#) (C function),  
[UV\\_UNKNOWN](#) (C macro),  
[uv\\_unref](#) (C function),  
[uv\\_update\\_time](#) (C function),  
[uv\\_uptime](#) (C function),  
[uv\\_utsname\\_t](#) (C type),  
[uv\\_version](#) (C function),  
[UV\\_VERSION\\_HEX](#) (C macro),  
[UV\\_VERSION\\_IS\\_RELEASE](#) (C macro),  
[UV\\_VERSION\\_MAJOR](#) (C macro),  
[UV\\_VERSION\\_MINOR](#) (C macro),  
[UV\\_VERSION\\_PATCH](#) (C macro),  
[uv\\_version\\_string](#) (C function),  
[UV\\_VERSION\\_SUFFIX](#) (C macro),  
[uv\\_walk](#) (C function),  
[uv\\_walk\\_cb](#) (C type),  
[uv\\_work\\_cb](#) (C type),  
[uv\\_work\\_t](#) (C type),  
[uv\\_work\\_t.loop](#) (C member),  
[uv\\_write](#) (C function),  
[uv\\_write2](#) (C function),  
[uv\\_write\\_cb](#) (C type),  
[uv\\_write\\_t](#) (C type),  
[uv\\_write\\_t.handle](#) (C member),  
[uv\\_write\\_t.send\\_handle](#) (C member),

[uv\\_handle\\_t.loop](#) (C member),  
[uv\\_handle\\_t.type](#) (C member),  
[UV\\_HANDLE\\_TYPE\\_MAP](#) (C macro),  
[uv\\_handle\\_type\\_name](#) (C function),  
[uv\\_has\\_ref](#) (C function),  
[uv\\_hrtime](#) (C function),  
[uv\\_idle\\_cb](#) (C type),  
[uv\\_idle\\_init](#) (C function),  
[uv\\_idle\\_start](#) (C function),  
[uv\\_idle\\_stop](#) (C function),  
[uv\\_idle\\_t](#) (C type),  
[uv\\_if\\_indextoiid](#) (C function),  
[uv\\_if\\_indextoname](#) (C function),