

glibc FILE buffer

20251025 Sat

Copyright (c) 2025 Young W. Lim.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

I used Microsoft Copilot to prepare this document



Microsoft Copilot

<https://copilot.microsoft.com>

Microsoft Copilot: Your AI companion

Microsoft Copilot is your companion to inform, entertain and inspire. Get advice, feedback and straightforward answers. Try Copilot now.

difference between glibc FILE buffer and buffer used to call system call
write read

The glibc FILE buffer is a user-space abstraction that adds efficiency and flexibility to I/O operations, while the buffer used in system calls like `read()` and `write()` is directly managed by the programmer and operates in kernel space.
by the caller of `read()`, `write()`

Key Differences Between glibc FILE Buffer and System Call Buffers

Feature	glibc FILE Buffer (<code>fread</code> , <code>fwrite</code>)	System Call Buffer (<code>read</code> , <code>write</code>)
Location	Managed in <u>user space</u> by the C standard library	Passed directly to the kernel via syscalls
Structure	Internally managed via the <code>FILE</code> struct	Raw memory buffer (e.g., <code>char buf[1024]</code>)
Buffering	Supports full , line , or no buffering	No automatic buffering; programmer controls it
Efficiency	More efficient for small or frequent I/O	Less efficient unless manually buffered
Flexibility	<u>Automatically</u> handles buffering and flushing	Requires <u>manual</u> handling of buffer size and flush
Functions Used	<code>fread()</code> , <code>fwrite()</code> , <code>fgetc()</code> , <code>fprintf()</code>	<code>read()</code> , <code>write()</code>
Error Handling	Uses <code>ferror()</code> and <code>feof()</code>	Uses return values and <code>errno</code>
Thread Safety	Often thread-safe (depending on implementation)	Not inherently thread-safe

Performance Implications

Using glibc's buffered I/O is generally faster for small or frequent reads/writes because it minimizes system calls. However, for large, bulk transfers or low-level control, direct system calls may be preferable.

Would you like a benchmark comparison or code samples showing how performance differs between the two approaches?

interaction between buffer in `fread(buffer, 1, 100, fp)`; and glibc FILE buffer

Great question — the interaction between the buffer you pass to `fread()` and the internal buffer managed by glibc's FILE structure is subtle but important. Here's how they work together:

user buffer

How `fread(buffer, 1, 100, fp)` Interacts with glibc's FILE Buffer

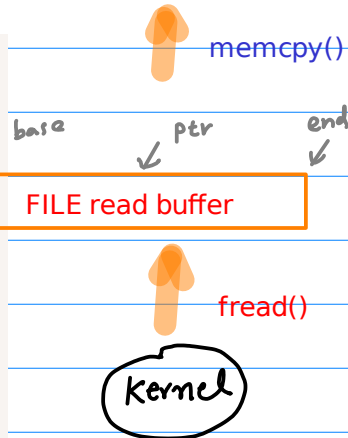
When you call `fread(buffer, 1, 100, fp)`, you're asking glibc to read 100 bytes from the stream `fp` into your user-provided buffer. Here's what happens under the hood:

1. glibc's Internal Buffer (FILE buffer)

- glibc maintains its own internal buffer for each FILE stream (e.g., `fp`).
- This buffer is defined by pointers like `_IO_read_base`, `_IO_read_ptr`, and `_IO_read_end`.
- It is used to reduce the number of expensive system calls (`read()`), by reading larger chunks from the file descriptor and serving smaller pieces to the user.

2. Your Buffer (`buffer`)

- This is the destination buffer you provide to `fread()`.
- glibc copies data from its internal buffer into your buffer.



[Kernel File Descriptor] <-- read() --> [glibc FILE internal buffer] <-- memcpy --> [Your buffer]

3. Interaction Flow

Text ^

Copy

[Kernel File Descriptor] <-- read() --> [glibc FILE internal buffer] <-- memcpy --> [Your buffer]

- If glibc's internal buffer already contains enough data (≥ 100 bytes), it simply copies from `_IO_read_ptr` to your buffer and advances the pointer.
- If not enough data is available, glibc performs a system call (`read()`) to refill its internal buffer from the file descriptor.
- Then it copies the requested amount to your buffer.

4. Efficiency

- This layered buffering system improves performance:
 - glibc reads large chunks (e.g., 8 KB) from the kernel.
 - Your program reads smaller chunks (e.g., 100 bytes) from glibc.
- This minimizes the number of system calls and speeds up I/O.

glibc's `FILE` internal buffer is a memory region managed by the C standard library to optimize I/O operations by reducing direct system calls. It temporarily stores data read from or written to a file stream.

🧠 How glibc's `FILE` Internal Buffer Works

When you use functions like `fread()`, `fwrite()`, `fgetc()`, or `fprintf()`, glibc doesn't interact with the kernel directly for every byte. Instead, it uses an internal buffer defined in the `FILE` structure to batch I/O operations.

🔧 Key Buffer Fields In `FILE` (glibc's `struct _IO_FILE`)

These fields are part of glibc's internal implementation and are defined in `libio/libio.h`:

Field	Purpose
<code>_IO_buf_base</code>	Start of the <u>allocated</u> buffer
<code>_IO_buf_end</code>	End of the buffer
<code>_IO_read_base</code>	Start of the <u>readable</u> portion
<code>_IO_read_ptr</code>	Current read position
<code>_IO_read_end</code>	End of the <u>readable</u> portion
<code>_IO_write_base</code>	Start of the <u>writable</u> portion
<code>_IO_write_ptr</code>	Current write position
<code>_IO_write_end</code>	End of the <u>writable</u> portion

These pointers help glibc manage how much data is available in the buffer and when to refill or flush it.

🔄 Buffer Lifecycle

- Allocation:** When a stream is opened (e.g., via `fopen()`), glibc allocates a buffer (typically 8 KB).
- Reading:**
 - Data is read from the kernel into the internal buffer via `read()`.
 - `fread()` then copies data from the internal buffer to your user buffer.
- Writing:**
 - Data is written to the internal buffer.
 - When full or flushed (`fflush()`), glibc writes it to the kernel via `write()`.

⚙ Buffering Modes

glibc supports three buffering modes:

- **Fully buffered:** Flushes only when full.
- **Line buffered:** Flushes on newline or full.
- **Unbuffered:** Writes immediately.

These are controlled via `setvbuf()` or inferred from the stream type (e.g., `stdout` is line-buffered if connected to a terminal).

🔧 Example: Reading with `fread()`

C ^

📄 Copy

```
FILE *fp = fopen("data.txt", "r");
char buffer[100];
fread(buffer, 1, 100, fp); // Copies from internal buffer to user buffer
```

Here, glibc may read 8 KB from the file into its internal buffer, then copy 100 bytes to `buffer`.

readable portion

In glibc's `FILE` structure, the **readable portion** refers to the segment of the internal buffer that contains data already fetched from the file (via a system call like `read()`) and is ready to be consumed by the program.

🧠 Breakdown of the Readable Portion

This portion is defined by three key pointers:

Field	Description
<code>_IO_read_base</code>	Points to the beginning of the readable data in the buffer
<code>_IO_read_ptr</code>	Points to the current read position (advances as data is consumed)
<code>_IO_read_end</code>	Points to the end of the readable data

Together, they define the range:

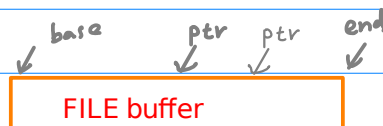
```
C ^  
[ _IO_read_ptr, _IO_read_end )
```

Copy

This is the active window of data that `fread()`, `fgetc()`, and similar functions use.

📦 How It Works in Practice

- Buffer Fill:** When the internal buffer is empty or exhausted, glibc calls `read()` to fill it from the file descriptor.
- Read Access:** Functions like `fread()` copy data from `_IO_read_ptr` up to `_IO_read_end` into your user buffer.
- Pointer Advance:** After reading, `_IO_read_ptr` is moved forward to reflect consumed data.
- Refill Trigger:** If `_IO_read_ptr == _IO_read_end`, the buffer is empty and needs refilling.



Example Flow

Suppose glibc reads 8 KB from a file into its internal buffer:

- `_IO_read_base` → start of buffer
- `_IO_read_ptr` → current position (e.g., 2 KB in)
- `_IO_read_end` → end of valid data (e.g., 8 KB)

If you call `fread(user_buf, 1, 100, fp)`, glibc copies 100 bytes from `_IO_read_ptr` and advances it by 100 bytes.

FILE vs. _IO_FILE

The difference between `FILE` and `_IO_FILE` in glibc is mostly about abstraction and implementation:



`FILE` vs `_IO_FILE` in glibc

Term	Description
<code>FILE</code>	Public-facing type defined in <code><stdio.h></code> used by C programs for stream I/O
<code>_IO_FILE</code>	<u>Internal structure</u> used by glibc to implement <code>FILE</code> functionality

Details

`FILE`

- Defined as a **typedef** in `<stdio.h>`:

```
typedef struct _IO_FILE FILE;
```

- Used in standard C functions like `fopen()`, `fread()`, `fprintf()`, etc.
- Acts as an opaque handle for file streams.

`_IO_FILE`

- Defined in glibc's internal header: `libio/libio.h`.
- Contains all the internal fields used to manage buffering, file descriptors, flags, and stream state.
- Includes fields like:
 - `_IO_read_ptr`, `_IO_read_end`
 - `_IO_write_ptr`, `_IO_write_end`
 - `_fileno`, `_flags`, `_lock`, etc.



Why the Separation?

- Encapsulation:** `FILE` hides implementation details from the user.
- Portability:** Programs written in C use `FILE` without worrying about glibc internals.
- Flexibility:** glibc can change `_IO_FILE` without breaking user code.

FILE structure

The `FILE` structure in glibc is a typedef for the internal `_IO_FILE` structure, which manages stream I/O operations like buffering, file descriptors, and read/write pointers.

Overview of `FILE` Structure in glibc

In glibc, the standard C `FILE` type is defined as:

```
C ^  
typedef struct _IO_FILE FILE;
```

 Copy

This means `FILE` is just an alias for `_IO_FILE`, which is the actual structure used internally by glibc to manage file streams. This structure is defined in `libio/libio.h` and is quite complex, designed to support buffered I/O, thread safety, and polymorphic stream behavior.

How It Works

- When you open a file with `fopen()`, glibc allocates a `_IO_FILE_plus` structure, which contains `_IO_FILE` and a pointer to a virtual function table (`_IO_jump_t`).
- This setup allows glibc to support different types of streams (e.g., files, sockets, memory buffers) using the same interface.
- The internal buffer is used to optimize I/O by reducing system calls.

Key Fields in `_IO_FILE`

Field	Description
<code>_flags</code>	Status flags (e.g., EOF, error, buffering mode)
<code>_IO_read_ptr</code>	Current read position in the buffer
<code>_IO_read_end</code>	End of readable data
<code>_IO_read_base</code>	Start of readable data
<code>_IO_write_base</code>	Start of writable buffer
<code>_IO_write_ptr</code>	Current write position
<code>_IO_write_end</code>	End of writable buffer
<code>_IO_buf_base</code>	Start of the allocated buffer
<code>_IO_buf_end</code>	End of the allocated buffer
<code>_fileno</code>	File descriptor associated with the stream
<code>_lock</code>	Mutex for thread safety
<code>_offset</code>	Current file offset
<code>_IO_save_base</code>	Base of backup area (used for ungetc and similar operations)
<code>_IO_backup_base</code>	Backup pointer for ungetc
<code>_IO_save_end</code>	End of saved area
<code>_IO_marker</code>	Linked list of position markers (used for stream positioning)
<code>_IO_chain</code>	Pointer to next stream in chain (used internally by glibc)
<code>_IO_jump_t *vtable</code>	Virtual function table for polymorphic stream operations

_iobuf

`_iobuf` is a structure used in Microsoft's C runtime library (MSVC) to represent file streams, similar to how glibc uses `_IO_FILE` in Linux systems. It's the internal implementation behind the standard `FILE` type in Windows-based C programs.



`_iobuf` Structure in MSVC

In Microsoft's implementation, `FILE` is defined as:

```
typedef struct _iobuf {
    char* _ptr;      // Current position in buffer
    int _cnt;        // Remaining characters in buffer
    char* _base;     // Pointer to the buffer
    int _flag;        // Flags (e.g., read/write mode, error, EOF)
    int _file;        // File descriptor
    int _charbuf;     // Type of buffering
    int _bufsiz;     // Size of buffer
    char* _tmpfname; // Temporary filename (if any)
} FILE;
```

Field Breakdown

Field	Purpose
<code>_ptr</code>	Points to the current position in the buffer
<code>_cnt</code>	Number of characters left to read/write
<code>_base</code>	Start of the buffer
<code>_flag</code>	Status flags (e.g., <code>_IOREAD</code> , <code>_IOWRT</code> , <code>_IOEOF</code> , <code>_IOERR</code>)
<code>_file</code>	File descriptor used in system calls
<code>_charbuf</code>	Indicates buffering mode (e.g., full, line, none)
<code>_bufsiz</code>	Size of the buffer
<code>_tmpfname</code>	Temporary filename (used for temp files)



Summary

- `_iobuf` is the Windows/MSVC equivalent of glibc's `_IO_FILE`.
- It underpins the standard `FILE` type used in functions like `fopen()`, `fread()`, `fwrite()`, etc.
- While the structure is similar in concept to glibc's version, the field names and layout differ.

in glibc, the FILE internal buffer is used to call read /write or additional buffer is used?

In glibc, the `FILE` structure's internal buffer is directly used for `read` and `write` operations — **no additional buffer is allocated** unless explicitly requested by the programmer.

How It Works

When you use standard I/O functions like `fread()`, `fwrite()`, `fgetc()`, or `fprintf()`, glibc performs buffering through the `FILE` object:

- The `FILE` structure contains an internal buffer (e.g., `_IO_buf_base`, `_IO_buf_end`)
- This buffer is used to **minimize system calls** by batching reads and writes
- The buffer is filled using `read()` and flushed using `write()` at the kernel level

in glibc, fopen

In glibc, the `fopen()` function is used to **open a file stream** and return a pointer to a `FILE` object, which is then used for buffered I/O operations like `fread()`, `fwrite()`, `fprintf()`, and `fscanf()`.

Function Prototype

C ^

 Copy

```
#include <stdio.h>
```

```
FILE *fopen(const char *pathname, const char *mode);
```

- `pathname` : path to the file
- `mode` : access mode (`"r"`, `"w"`, `"a"`, `"r+"`, etc.)

What Happens Internally in glibc

When you call `fopen()`:

1. glibc allocates a `FILE` structure
2. It opens the file using the system call `open()`
3. It sets up an internal buffer for efficient I/O
4. It returns a pointer to the `FILE` object

This buffer is used for **read/write operations**, minimizing direct system calls.

Buffering Behavior

- Controlled by `setvbuf()` or `setbuf()`
- Default buffering:
 - **Files** → fully buffered
 - **Terminal (stdout)** → line buffered
 - **stderr** → unbuffered



how `fopen()` interacts with low-level system calls or how to customize buffering for performance?

Great question — let's break this into two parts: how `fopen()` interacts with low-level system calls, and how you can customize buffering for performance.

Part 1: How `fopen()` Interacts with Low-Level System Calls


When you call `fopen()` in glibc:

1. `fopen()` **allocates a `FILE` structure** in user space.
2. It **calls the `open()` system call** under the hood to get a file descriptor (an integer handle to the file).
3. It **associates the file descriptor with the `FILE` stream**, initializing internal buffers and setting up the appropriate read/write modes.
4. All subsequent I/O (e.g., `fread()`, `fwrite()`, `fgetc()`) uses **buffered I/O**, which means:
 - Data is read/written in chunks using `read()` / `write()` system calls.
 - This reduces the number of expensive kernel transitions.



Example: `fread()` Flow

```
fread(buf, 1, 100, fp);
```

 Copy

- If the internal buffer is empty, glibc calls `read(fd, internal_buffer, BUFSIZ)`
- Then it copies up to 100 bytes from the internal buffer to `buf`
- If the buffer still has data, no system call is made on the next `fread()`



Part 2: Customizing Buffering for Performance

You can control how glibc buffers I/O using:

1. `setvbuf()`

C ^

Copy

```
int setvbuf(FILE *stream, char *buf, int mode, size_t size);
```

- `mode` can be:
 - `_IOFBF` – Full buffering
 - `_IOLBF` – Line buffering
 - `_IONBF` – No buffering
- `buf` can be:
 - `NULL` → glibc allocates its own buffer
 - Custom buffer → you provide memory
- `size` sets the buffer size

Example: Custom 8KB Buffer

```
C ^  
char mybuf[8192];  
FILE *fp = fopen("data.txt", "r");  
setvbuf(fp, mybuf, _IOFBF, sizeof(mybuf));
```

 Copy

This reduces the number of `read()` calls and can improve performance for large sequential reads.

Important Notes

- Always call `setvbuf()` **immediately after** `fopen()`, before any I/O
- For performance-critical applications (e.g., loggers, large file processors), tuning buffer size can yield noticeable gains

in glibc, what is the exact definition of stream

In glibc, a **stream** refers to an abstraction for handling input and output — typically represented by a `FILE *` object. This stream encapsulates all the internal state needed to manage buffered I/O operations on files, sockets, pipes, or other data sources.

Exact Definition in glibc

According to the GNU C Library documentation:

“A `FILE` object holds all of the internal state information about the connection to the associated file, including such things as the file position indicator and buffering information.” GNU Wget

Internally, glibc implements `FILE` as a complex structure called `_IO_FILE`, which is part of the **libio** subsystem. When you use `fopen()`, glibc allocates a `struct _IO_FILE_plus`, which includes:

- A `struct _IO_FILE` for stream state
- A pointer to a virtual function table (`_IO_jump_t`) for operations like `fread`, `fwrite`, etc.

HackMD


What a Stream Tracks

A glibc stream (`FILE *`) manages:

- File descriptor
- Buffer pointers and size
- Current position in the buffer
- Error and EOF flags
- Locking for thread safety
- Function pointers for I/O operations

Example

```
C ^  
FILE *fp = fopen("data.txt", "r");  
fgetc(fp); // Reads from the stream  
fclose(fp); // Closes the stream
```

 Copy

Here, `fp` is a stream that wraps the file descriptor and handles buffering, error tracking, and more.

Related Structures

- `struct _IO_FILE` → core stream state
- `struct _IO_FILE_plus` → adds function dispatch
- `struct _IO_jump_t` → virtual function table for stream operations

You can explore these in glibc source under `libio/libio.h`.

what v stands for, in setvbuf

Great question! In `setvbuf()`, the `v` stands for "variable" — as in "set variable buffer".

Here's the breakdown:

- `setvbuf` = set variable buffer
- It allows you to set a custom buffer (`buf`) and choose the buffering mode (`_IOFBF`, `_IOLBF`, `_IONBF`)
- This is more flexible than `setbuf()`, which only lets you assign a buffer without specifying mode or size

Related Functions

Function	Meaning	Flexibility
<code>setbuf()</code>	Set buffer (fixed behavior)	✗ Limited
<code>setvbuf()</code>	Set variable buffer	✓ Full control

So the `v` is all about giving you **variable control** over how buffering works.

In C (and Linux), `setvbuf()` is a standard library function used to control **buffering behavior** for a `FILE*` stream. It allows you to choose the buffer mode and size, which can affect performance and I/O timing.

Function Prototype

C ^

 Copy

```
#include <stdio.h>

int setvbuf(FILE *stream, char *buf, int mode, size_t size);
```

Parameters:

- `stream` : the `FILE*` stream (e.g., from `fopen()`)
- `buf` : pointer to a buffer (or `NULL` to let the system allocate one)
- `mode` : buffering mode (`_IOFBF`, `_IOLBF`, `_IONBF`)
- `size` : size of the buffer



Buffering Modes

Mode	Description
<code>_IOFBF</code>	Full buffering (waits until buffer full)
<code>_IOLBF</code>	Line buffering (flushes on newline)
<code>_IONBF</code>	No buffering (writes immediately)



Example

```
C ^ Copy  
  
FILE *fp = fopen("output.txt", "w");  
char buffer[1024];  
setvbuf(fp, buffer, _IOFBF, sizeof(buffer));
```

This sets up a 1 KB fully buffered stream for `fp`.



Notes

- Must be called **before any I/O** on the stream — otherwise it has no effect.
- Useful for tuning performance in high-throughput or real-time systems.
- If `buf == NULL`, glibc will allocate a buffer internally.