

HOW LINUX PIPES WORK UNDER THE HOOD

Piping is one of the core concepts of Linux & Unix based operating systems. Pipes allow you to chain together commands in a very elegant way, passing output from one program to the input of another to get a desired end result.

Here's a simple example of piping:

```
ls -la | sort | less
```

The above command gets a listing of the current directory using `ls`, sorts it alphabetically using the `sort` utility, and then paginates it for easier reading using `less`.

I'm not going to go into depth about pipes as I'll assume you already know what they are (at least in the context of bash pipes). Instead, I'm going to show how pipes are implemented under the hood.

How Pipes Are Implemented

Before I explain how bash does pipes, I'll explain how the kernel implements pipes (at a high level).

- Linux has a VFS (virtual file system) module called pipefs, that gets mounted in kernel space during boot
- pipefs is mounted alongside the root file system (`/`), not in it (pipe's root is `pipe:`)
- pipefs cannot be directly examined by the user unlike most file systems
- The entry point to pipefs is the `pipe(2)` syscall
- The `pipe(2)` syscall is used by shells and other programs to implement piping, and just creates a new file in pipefs, returning two file descriptors (one for the read end, opening using `O_RDONLY`, and one for the write end, opened using `O_WRONLY`)
- pipefs is stored as an in-memory file system

Pipe I/O, buffering, and capacity

A pipe has a limited capacity in Linux. When the pipe is full, a `write(2)` will block (or fail if the `O_NONBLOCK` flag is set). Different implementations of pipes have different limits, so applications shouldn't rely on a pipe having a particular size. Applications should be designed to consume data as soon as it is available so the writing process doesn't block. That said, knowing the pipe size is useful. Since

Linux 2.6.35, the default pipe capacity is 65,536 bytes (it used to be the size of the page file, e.g. 4096 bytes in i386 architectures).

When a process attempts to read from an empty pipe, `read(2)` will block until data is available in the pipe. If all file descriptors pointing to the write end of the pipe have been closed, reading from the pipe will return EOF (`read(2)` will return 0).

If a process attempts to write to a full pipe, `write(2)` will block until enough data has been read from the pipe to allow the write call to succeed. If all file descriptors pointing to the read end of the pipe have been closed, writing to the pipe will raise the SIGPIPE signal. If this signal is ignored, `write(2)` fails with the error `EPIPE`.

All of this is important when understanding pipe performance. If a process A is writing data at roughly the same speed as process B is reading it, pipes work very well and are highly performance. An imbalance here can cause performance problems. See the next section for more information/examples.

How Shells Do Piping

Before continuing, you should be aware of [how Linux creates new processes](#).

Shells implement piping in a manner very similar to [how they implement redirection](#). Basically, the parent process calls `pipe(2)` once for each two processes

that get piped together. In the example above, bash would need to call `pipe(2)` twice to create two pipes, one for piping `ls` to `sort`, and one to pipe `sort` to `less`. Then, bash forks itself once for each process (3 times for our example). Each child will run one command. However, before the children run their commands, they will overwrite one of `stdin` or `stdout` (or both). In our above example, it will work like this:

- bash will create two pipes, one to pipe `ls` to `sort`, and one to pipe `sort` to `less`
- bash will fork itself 3 times (1 parent process and 3 children, for each command)
- child 1 (`ls`) will set its `stdout` file descriptor to the write end of pipe A
- child 2 (`sort`) will set its `stdin` file descriptor to the read end of pipe A (to read input from `ls`)
- child 2 (`sort`) will set its `stdout` file descriptor to the write end of pipe B
- child 3 (`less`) will set its `stdin` file descriptor to the read end of pipe B (to read input from `sort`)
- each child will run their commands

The kernel will automatically schedule processes so they roughly run in parallel. If child 1 writes too much to pipe A before child 2 has read it, child 2 will block for a while until child 2 has had time to read from the pipe. This normally allows for

very high levels of efficiency as one process doesn't have to wait for the other to complete to start processing data. Another reason for this is that pipes have a limited size (normally the size of a single page of memory).

Pipe Example Code

Here is a C example of how a program like bash might implement piping. My example is pretty simple, and accepts two arguments: a directory and a string to search for. It will run `ls -la` to get the contents of the directory, and pipe them to `grep` to search for the string.

```
1. #include <unistd.h>
2. #include <stdio.h>
3. #include <errno.h>
4. #include <fcntl.h>
5.
6. #define READ_END 0
7. #define WRITE_END 1
8.
9. int main(int argc, char *argv[])
10. {
11.     int pid, pid_ls, pid_grep;
12.     int pipefd[2];
13.
14.     // Syntax: test . filename
15.     if (argc < 3) {
16.         fprintf(stderr, "Please specify the directory to search and the filename to search for\n");
17.         return -1;
18.     }
19.
20.     fprintf(stdout, "parent: Grepping %s for %s\n", argv[1], argv[2]);
21.
```

```
22. // Create an unnamed pipe
23. if (pipe(pipefd) == -1) {
24.     fprintf(stderr, "parent: Failed to create pipe\n");
25.     return -1;
26. }
27.
28. // Fork a process to run grep
29. pid_grep = fork();
30.
31. if (pid_grep == -1) {
32.     fprintf(stderr, "parent: Could not fork process to run grep\n");
33.     return -1;
34. } else if (pid_grep == 0) {
35.     fprintf(stdout, "child: grep child will now run\n");
36.
37.     // Set fd[0] (stdin) to the read end of the pipe
38.     if (dup2(pipefd[READ_END], STDIN_FILENO) == -1) {
39.         fprintf(stderr, "child: grep dup2 failed\n");
40.         return -1;
41.     }
42.
43.     // Close the pipe now that we've duplicated it
44.     close(pipefd[READ_END]);
45.     close(pipefd[WRITE_END]);
46.
47.     // Setup the arguments/environment to call
48.     char *new_argv[] = { "/bin/grep", argv[2], 0 };
49.     char *envp[] = { "HOME=/", "PATH=/bin:/usr/bin", "USER=brandon", 0 };
50.
51.     // Call execve(2) which will replace the executable image of this
52.     // process
53.     execve(new_argv[0], &new_argv[0], envp);
54.
55.     // Execution will never continue in this process unless execve returns
56.     // because of an error
57.     fprintf(stderr, "child: Oops, grep failed!\n");
58.     return -1;
59. }
60.
```

```

61. // Fork a process to run ls
62. pid_ls = fork();
63.
64. if (pid_ls == -1) {
65.     fprintf(stderr, "parent: Could not fork process to run ls\n");
66.     return -1;
67. } else if (pid_ls == 0) {
68.     fprintf(stdout, "child: ls child will now run\n");
69.     fprintf(stdout, "-----\n");
70.
71.     // Set fd[1] (stdout) to the write end of the pipe
72.     if (dup2(pipefd[WRITE_END], STDOUT_FILENO) == -1) {
73.         fprintf(stderr, "ls dup2 failed\n");
74.         return -1;
75.     }
76.
77.     // Close the pipe now that we've duplicated it
78.     close(pipefd[READ_END]);
79.     close(pipefd[WRITE_END]);
80.
81.     // Setup the arguments/environment to call
82.     char *new_argv[] = { "/bin/ls", "-la", argv[1], 0 };
83.     char *envp[] = { "HOME=", "PATH=/bin:/usr/bin", "USER=brandon", 0 };
84.
85.     // Call execve(2) which will replace the executable image of this
86.     // process
87.     execve(new_argv[0], &new_argv[0], envp);
88.
89.     // Execution will never continue in this process unless execve returns
90.     // because of an error
91.     fprintf(stderr, "child: Oops, ls failed!\n");
92.     return -1;
93. }
94.
95. // Parent doesn't need the pipes
96. close(pipefd[READ_END]);
97. close(pipefd[WRITE_END]);
98.
99. fprintf(stdout, "parent: Parent will now wait for children to finish execution\n");

```

```
100.
101.     // Wait for all children to finish
102.     while (wait(NULL) > 0);
103.
104.     fprintf(stdout, "-----\n");
105.     fprintf(stdout, "parent: Children has finished execution, parent is done\n");
106.
107.     return 0;
108. }
```

I've commented it thoroughly, so hopefully it makes sense.

Named vs Unnamed Pipes

In the above examples, we've been using unnamed/anonymous pipes. These pipes are temporary, and are discarded once your program finishes or all of their file descriptors are closed. They are the most common type of pipe.

Named pipes, also known as FIFOs (for first in, first out), get created as a named file on your hard disk. They allow multiple unrelated programs to open and use them. You can have multiple writers quite easily, with one reader, for a very simplistic client-server type design. For example, nagios does this, with the master process reading a named pipe, and every child process writing commands to the named pipe.

Named pipes are creating using the mkfifo command or syscall. Example:

```
mkfifo ~/test_pipe
```

Other than their creation, they work pretty much the same as unnamed pipes. Once you create them, you can open them using `open(2)`. You must open the read end using `O_RDONLY` or the write end using `O_WRONLY`. Most operating systems implement unidirectional pipes, so you can't open them in both read/write mode.

FIFOs are often used as a unidirectional IPC technique, for a system with multiple processes. A multithreaded application may also use named or unnamed pipes, as well as other IPC techniques such as shared memory segments.

FIFOs are created as a single inode, with the property `i_pipe` set as a reference to the actual pipe. While the name may exist on your filesystem, pipes don't cause I/O to the underlying device, as once the inode is read, FIFOs behave like unnamed pipes and operate in-memory.

Source: <http://brandonwamboldt.ca/how-linux-pipes-work-under-the-hood-1518/>