



Algorithms in the Real World

ELF Files



A “file format” is just a specification that gets distributed and agreed upon.

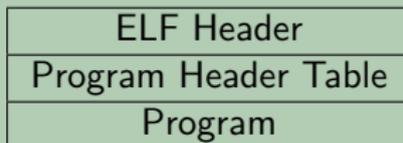
Your computer first figures out which type of file it is using the first few bytes of the file. This is called a “magic number”:

Magic Numbers

- PDF files start with 25 50 44 46.
- JPEG image files begin with FF D8.
- ELF executables start with 7F E L F.

Then, it runs a program that understands the specification.

Pieces



e_ident Field

Field Name	Length	Description
magic number	4 bytes	Use 0x7F 'E' 'L' 'F'
class	1 byte	Use 1 for 32-bit and 2 for 64-bit
data	1 byte	Use 1 for little and 2 for big
version	1 byte	Use 1
pad	8 bytes	Use 0
size	1 byte	Use 16

Minimal ELF files consist of three pieces: the ELF header, the program header, and the program. We'll construct our file as follows:

- 1 Write our program.
- 2 Fill in the program header according to the specification.
- 3 Fill in as much of the header as possible according to the specification.

For the purposes of this exercise, we don't really care what our program does. So, we'll go ahead and exit with an exit code of 66. That is, we'll write the following assembly code:

```
1 BITS 32
2 mov ebx, 0x42 # Set the return value to 66
3 mov eax, 0x1 # Call the 'exit' syscall
4 int 0x80     # ...via interrupt 0x80
```

Compiling it with `nasm` gives us: OUTPUT

```
>> bb 42 00 00 00 b8 01 00 00 00 cd 80
```

Minimal ELF files consist of three pieces: the ELF header, the program header, and the program. We'll construct our file as follows:

- 1 Write our program.
- 2 Fill in the program header according to the specification.
- 3 Fill in as much of the header as possible according to the specification.

Our Program

```
bb 42 00 00 00 b8 01 00 00 00 cd 80
```

Steps 2 and 3 are basically just 'follow the specification'. We'll do step 2 together, and you can try step 3 on your own.