



Flare-On Challenge 8 Solution

By Eamon Walsh

Challenge 3: Antioch



Challenge Prompt

To solve this challenge, you'll need to ...AAARGH

Solution

The challenge zip package contains a tar archive `antioch.tar`. Inside the archive is the following file structure:

```
manifest.json
repositories
09e6fff53d6496d170aaa9bc88bd39e17c8e5c13ee9066935b089ab0312635ef/
  json
  layer.tar
  VERSION
1c5d28d6564aed0316526e8bb2d79a436b45530d2493967c8083fea2b2e518ce/
  json
  layer.tar
  VERSION
.../
```

The top level contains `manifest.json`, `repositories`, and a JSON file named with a hexadecimal ID. There are 31 additional folders named with a hexadecimal ID and containing `json`, `layer.tar`, and `VERSION` files. Examining the names and contents of the files and doing some basic Internet searches for those names, we identify this archive as an exported Docker image.

A Docker image is a virtual filesystem and launch configuration that form the basis for a container, an isolated execution environment in which one or more processes can run, typically with a complete set of libraries and other dependencies. The flexibility of these containers has made Docker images a popular choice for the packaging and distribution applications.

Can the image be loaded? From within a Linux analysis VM with Docker installed:

```
$ docker image load -i antioch.tar
d26c760acd6e: Loading layer 14.85kB/14.85kB
Loaded image: antioch:latest
$ docker image ls -a
REPOSITORY TAG IMAGE ID CREATED SIZE
antioch latest a13ffcf46cf4 6 weeks ago 13kB
```

The image loads successfully, which means we can use Docker tools to examine and run the image. A Docker image consists of a sequence of layers, each of which makes a change to the image. Most layers add files to the virtual filesystem, but layers can also remove files or adjust the launch configuration for containers generated from the image. Using the `history` command, the layers within the image can be viewed:

```
$ docker history a13ffcf46cf4
IMAGE CREATED BY SIZE
a13ffcf46cf4 /bin/sh -c #(nop) CMD ["/AntiochOS"] 0B
<missing> /bin/sh -c #(nop) ADD file:fae1674... in / 13kB
```

The image is quite simple, consisting of only two layers. The first layer adds a file "AntiochOS" to the root directory of the (initially empty) virtual filesystem and the second layer configures the container to run that file on startup. Many Docker images have a large virtual filesystem with a complete operating system runtime, but it's not uncommon to find an image containing statically linked executables that only require the kernel interfaces to run. For example, programs written in Go are typically compiled statically and packaged this way. However, the small size (13kB) of our file suggests that it is not a Go executable.

Let's run the container and see what happens:

```
$ docker run -it --rm --name antioch_test a13ffcf46cf4
AntiochOS, version 1.32 (build 1975)
Type help for help
> help
Available commands:
help: print this help
...AAARGH
>
```

The program presents a CLI interface, although the “help” command is not terribly helpful, and other guessed commands do not produce a response. To analyze the binary, we'll need to extract it from the image. One way to do this is to copy it out of a running container:

```
$ docker cp antioch_test:/AntiochOS .
$ ls
AntiochOS
$
```

Another way is to find its layer in `antioch.tar` and extract it from the `layer.tar` file in that folder. This method exposes a detail that sharp-eyed readers might have already picked up on: There are extra layers in `antioch.tar` that aren't listed in the image history.

Examining the `layer.tar` for these extra layers, we find series of files named with a single letter: `a.dat`, `b.dat`, etc. all with size 4096 and seemingly random contents, some with zero contents. Another strange detail is found in the json descriptor files for these layers. Most of the JSON consists of boilerplate, with many null and empty values, but the “author” field seems to have significant content:

```
$ cat */json | jq '.author' | sort
"A Famous Historian"
"Black Knight"
"Bridge Keeper"
"Brother Maynard"
"Chicken of Bristol"
"Dennis the Peasant"
...
```

We'll keep this in mind for later. Returning to the AntiochOS binary, we note that it is a statically linked, stripped ELF executable. Running strings on it yields nothing useful, which is unexpected since we know that it prints strings as part of its CLI. This is a strong clue that string obfuscation techniques are present in the program. The next step is to disassemble the file in IDA Pro or Ghidra and/or load it into a debugger of choice.

Near the beginning of the entry point function start (0x401834) is the first obfuscated string, the banner/version string. The function at 0x4013E0 uses XOR with the constant 0x8B to deobfuscate the buffer at 0x4040C0. The first byte of the buffer starts as zero and is used to check if the buffer has already been deobfuscated. The next bytes in the buffer are the string itself. The functions 0x4010A0, 0x4010E0, 0x401260, and 0x4012A0 are obfuscated strings using the same scheme. Another obfuscated string is seen immediately afterwards at the function 0x4012E0 (“Type help for help”). This is a stack string, a type of obfuscated string that avoids the use of a separate buffer by placing the string bytes in the program code as immediate moves. The functions 0x401000, 0x401120, 0x401180, 0x4011E0, 0x4012E0, 0x401340, 0x401360, 0x401380, and 0x4013B0 are stack strings as well.

The rest of the entry point function is the loop that handles command entry. By examining the stack strings in the loop, we can determine the accepted commands: `help`, `quit`, `consult`, and `approach`. The next step is to determine what the `consult` and `approach` commands do. The function 0x401460 is the handler for the `consult` command. It is somewhat long but can be divided into several sections after review. The first part of the function sets up local variables, prints out a string, and calls a 2 second sleep. The second part of the `consult` function is the loop at 0x4014EA. The loop variable is stored in `r14` and goes from ‘a’ to ‘z’ within the loop. Each iteration of the loop

opens the file `X.dat` where `X` is the loop variable and reads the contents (the functions `0x4019D0`, `0x401A10`, and `0x401A30` are `open`, `read`, and `close` syscalls respectively). XOR is used to combine all the file contents into a stack buffer whose address is stored in `r12`. The third part of the `consult` function at `0x401544` sets up a 256-character lookup table using the stack string at `0x401000` with the remainder of the table filled with the character `'.'` (the XMM instructions at `0x4015AA` are GCC's `builtin_memset` implementation).

The final part of the `consult` function is the loop at `0x401560`, which looks up each byte of the combined file contents in the lookup table, except for every 16th byte which is set to a newline character `0x0A` instead. Finally, the converted buffer is printed out. The characters in the lookup table are ASCII characters typically seen in ASCII-art "figlet" style letters, suggesting that the result is intended to be a readable message.

Running the `consult` command in a container launched from the image prints nothing but the letter `V`, since there are no `.dat` files present in the root directory, a zero-initialized buffer ends up being decoded. The implementation of the `consult` command indicates that the data in the `.dat` files is important and a certain combination of files is needed in the root directory to be decoded. There are many such files in the unused layers in `antioch.tar`, which we noted earlier. We could try the files from each individual layer, but this only produces garbled results.

Going back to the concept of layers in a Docker image, we note that each layer applies a change to the virtual filesystem of the image. If two layers both add a certain file to the filesystem, the file that is present in the final image will be the one in whichever layer was applied last. Files from earlier layers will "show through" in the final image only if no later layer replaces them. In what order are the unused layers in `antioch.tar` intended to be applied? We can't determine this from the JSON metadata, so we'll need to keep analyzing the binary and look at the `approach` command.

The function `0x401640` is the handler for the `approach` command. From the obfuscated string functions called within it, we can see that it asks a series of three classic questions. The first question asks for a name which is passed to the function `0x401B50`. The return value of that function is then used in the loop at `0x4016DA`, which compares it to every third 4-byte value in the data buffer `0x402000` until a match is found or 30 checks are made. The function `0x401B50` is a hash function, the third and final string obfuscation technique in the program. Specifically, the function computes the CRC32 of a string, which can be determined by entering a known string at the prompt, observing the function return value in a debugger, and comparing it to a CRC32 computed elsewhere, such as an online tool (one complicating factor is that the newline entered at the prompt is included in the CRC32 calculation). Another clue is the table at `0x402260`, which contains precomputed constants for the polynomial used in the CRC algorithm.

What 30 names correspond to the CRC32 values in the table at `0x402000`? Going back to the suspicious "author" values in the JSON metadata for the unused layers, we can try one of them and verify that it works and triggers the success path. Further analysis of the `approach` function shows that the answer to the second question can be any nonempty string, and that the answer to the third question is hashed to a CRC32 value which is compared to the 4-byte word in the `0x402000` table immediately following the CRC32 value for the name.

The layer metadata doesn't contain any obvious color names that might answer the third question, but we note through analysis of the `approach` function, that we can simply bypass the check and go to the success condition at `0x401771`. This can be done by modifying the program code or by replacing the color CRC32 values in `0x402000` table with a known CRC32 value and entering the corresponding string, for example:

```
> approach
Approach the Gorge of Eternal Peril!
What is your name? Bridge Keeper
What is your quest? a
What is your favorite color? a
Right. Off you go. #14
```

The number printed in the final message is the third four-byte value in the `0x402000` table entry, and we note that these values range from 1 to 30 indicating a possible layer ordering. Therefore each 12-byte entry in the `0x402000` table consists of the following three fields:

- CRC32 of the name
- CRC32 of the favorite color

- Possible layer order number

We can determine the layer order numbers for each layer in `antioch.tar` by entering each name manually or by computing the CRC32 values for all the names (remembering to add a trailing newline) and finding the matching value in the `0x402000` table. The resulting ordering is:

- 1 Miss Islington
- 2 Sir Bors
- 3 Tim the Enchanter
- 4 Dragon of Angnor
- 5 Brother Maynard
- 6 Sir Bedevere
- 7 Sir Robin
- 8 Zoot
- 9 Squire Concorde
- 10 Green Knight
- 11 Trojan Rabbit
- 12 Chicken of Bristol
- 13 Roger the Shrubber
- 14 Bridge Keeper
- 15 Sir Gawain
- 16 Legendary Black Beast of Argh
- 17 A Famous Historian
- 18 Sir Lancelot
- 19 Lady of the Lake
- 20 Rabbit of Caerbannog
- 21 Sir Not-Appearing-in-this-Film
- 22 Prince Herbert
- 23 King Arthur
- 24 Inspector End Of Film
- 25 Sir Ector
- 26 Squire Patsy
- 27 Dennis the Peasant
- 28 Dinky
- 29 Black Knight
- 30 Sir Gallahad

To produce the set of `.dat` files corresponding to this layer ordering, we can create an empty directory and extract each `layer.tar` into this directory in the given order (the `-C` flag to `tar` is helpful here). Running the AntiochOS program within this directory and then executing the `consult` command will reveal the flag in ASCII art. It is **"Five-Is-Right-Out@flare-on.com"**.

