

Broken, Abandoned, and Forgotten Code, Part 6

Author: Zachary Cutlip

Note: *It is assumed that the reader is debugging the processes described in this and the next several posts using emulation and IDA Pro. Those topics are outside the scope of this series and are covered in detail [here](#) and [here](#).*

In the [previous post](#), we switched gears and started looking at the web server for the Netgear R6200. That's because the HTTP daemon's code for upgrading the firmware is less broken and easier to analyze. We also analyzed a stock firmware image downloaded from Netgear to see how it is composed. Craig Heffner's binwalk identified three parts, a TRX header at offset 58, followed by a compressed Linux kernel, followed by a squashfs filesystem. All of those parts are well understood, which only leaves the first 58 bytes to analyze.

With the goal of recreating the header using a stock TRX header, Linux kernel, and filesystem, I described how we can use [Bowcaster](#) to create fake header data to aid in debugging. When we left off, I had started discussing httpd's `abCheckBoardID()` function at `0x0041C3D8`, which partially parses the firmware header. We identified a magic signature that should be at the firmware image's offset 0, as well as some sort of size field that should be at offset 4. We also discovered this header should be big endian encoded even though the target system is little endian.

In this part, we'll clarify the purpose of the size field as well as identify a checksum field. Identification of the checksum algorithm is tricky if you don't have an eye for that sort of thing (I do not). I'll show how to deal

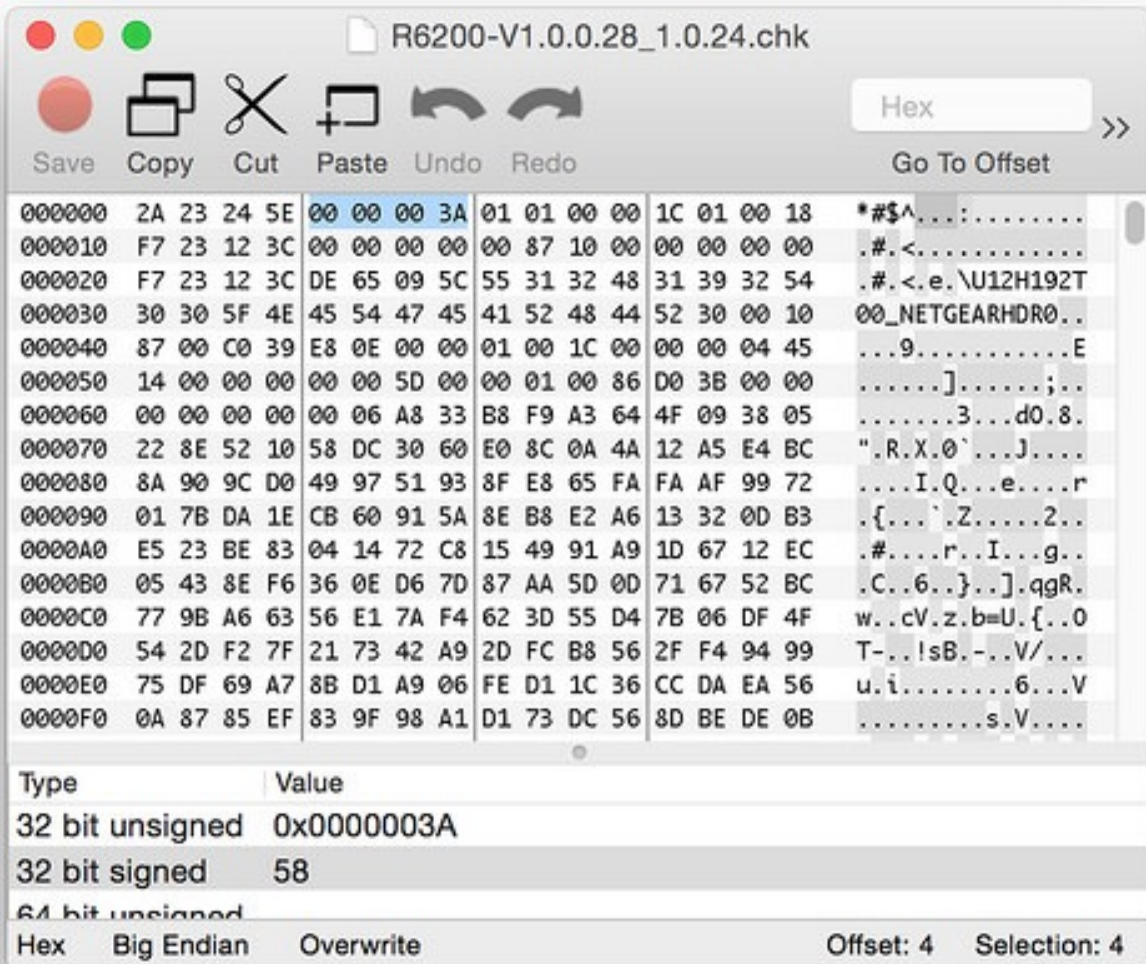
with that. By the end of this part, we will have identified four fields, accounting for 30 bytes of the 58-byte firmware header.

Updated Exploit Code

I last updated the exploit code for part 5, which added several Python modules to aid in reverse engineering and reconstructing a firmware image. In this part I've added a module to regenerate checksums found in the header (see below). Additionally, the `MysteryHeader` class populates a couple of new fields that we will cover this post. If you've previously cloned the repository, now would be a good time to do a pull. You can clone the git repo from:
https://github.com/zcutlip/broken_abandoned

Header Size

We know the field at offset 4 is a size field of some sort because it's used as the size for a `memcpy()` operation[1]. Let's take a look at a stock firmware image to see what value is in that field. It might correlate to something obvious.



Above, we see the stock value is 0x0000003A, or 58 in decimal. Since 58 is also the amount of unidentified data before the TRX header, it's a safe bet this field is the overall size of this unidentified header. It's also a safe bet that this header is variable in size. The TRX header, whose size is fixed, does not have a size field for the header alone, only for the header plus data.

```

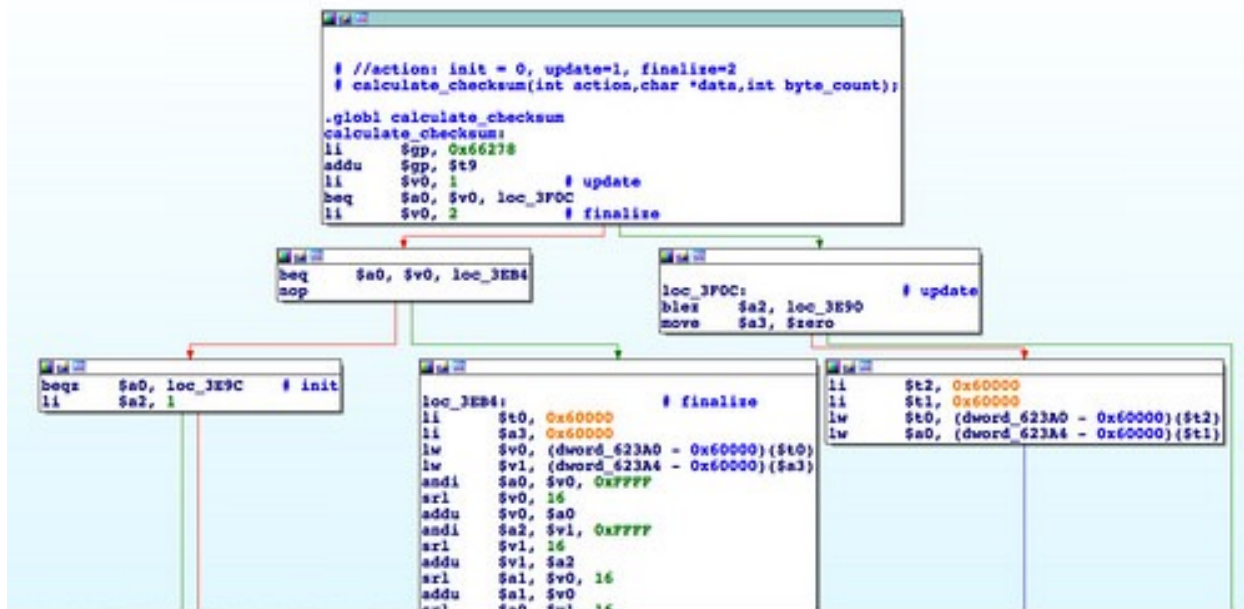
0041C528      addiu   $s2, $sp, 0xA0+var_88
0041C52C      la     $t9, memset
0041C530      move   $a0, $s2 # s
0041C534      move   $a1, $zero # c
0041C538      jalr   $t9 ; memset
0041C53C      li     $a2, 0x64 # 'd' # n
0041C540      lw     $gp, 0xA0+var_90($sp)
0041C544      move   $a1, $a1 # src
0041C548      la     $t9, memcpy
0041C54C      move   $a0, $s2 # dest
0041C550      jalr   $t9 ; memcpy
0041C554      move   $a2, $s0 # n
0041C558      lw     $gp, 0xA0+var_90($sp)
0041C55C      move   $a0, $zero
0041C560      la     $t9, calculate_checksum
0041C564      move   $a1, $zero # initialize the checksum
0041C568      move   $a2, $zero
0041C56C      jalr   $t9 ; calculate_checksum
0041C570      move   $a1, $t9
0041C574      lw     $gp, 0xA0+var_90($sp)
0041C578      move   $a1, $s2 # portion of firmware memcpy()ed above.
                                # This should be the 58 byte header
                                # with checksum field zeroed out.
0041C578      la     $t9, calculate_checksum
0041C57C      move   $a2, $s0 # size value used in memcpy()
                                # This is at offset 4 in header.
0041C580      move   $t9, $a1
0041C584      jalr   $t9
0041C588      li     $a0, 1
0041C590      lw     $gp, 0xA0+var_90($sp)
0041C594      li     $a0, 2
0041C598      la     $t9, calculate_checksum
0041C59C      move   $a1, $zero # finalize the checksum.
0041C5A0      move   $t9, $a1
0041C5A4      jalr   $t9
0041C5A8      move   $a2, $zero
0041C5AC      addu   $v1, $s3, $s4
0041C5B0      addu   $v1, $s6

```

Checksumming the firmware header.

Checksum Fun

From `abCheckBoardID()` there are several calls to the `calculate_checksum()` function. This is an imported symbol and is not in the `httpd` binary itself. Strings analysis of libraries on the R6200's filesystem reveals that this function is in the shared library `libacos_shared.so`. We can disassemble this binary and analyze the function.



Disassembly of calculate_checksum().

There's no need to completely reverse engineer this function. Sure, it would be convenient to know what checksum algorithm this is[2] and if there was a built-in python module to use. All we really need, however, is code that calculates the same values this function does. It's easier in this case to just reimplement the algorithm. I duplicated this function one-for-one, where each line of MIPS disassembly became a line of Python. It's a small function, so it didn't take long to do. That module is included in this week's update to the git repo.

```

26
27 ▼ def _update(self, data):
28     size=len(data)
29     t0=self.dword_623A0
30     a0=self.dword_623A4
31     a2=size
32     a3=0
33 ▼ while a3 != a2:
34     v1=ord(data[a3])
35     a3+=1
36     a0=(a0+v1) & 0xffffffff
37     t0=(t0+a0) & 0xffffffff
38
39     self.dword_623A0=t0
40     self.dword_623A4=a0
41
42     return 1
43
44 ▼ def _finalize(self):
45     v0=self.dword_623A0
46     v1=self.dword_623A4
47
48     a0=(v0 & 0xffff)
49     v0=(v0>>16 )
50     v0=(v0+a0) & 0xffffffff

```

Python code fragment that looks suspiciously like IDA Pro disassembly.

A checksum is calculated across the first 58 bytes of the header. Then at 0x0041C5BC the checksum gets compared to 0x41623241, a value extracted from the firmware data. Using Bowcaster's `find_offset()`, it is revealed that offset 36 of the firmware header should contain the checksum of the header itself. We'll need to calculate that value for the header and insert it at this location. In `abCheckBoardID()` the checksum field is zeroed out before the value is calculated. We should do the same before calculating our own. The updated code in the git repository performs this operation.

Board ID String

With the header checksum in place, we can move forward to the next few basic blocks. A few checks are performed to verify the "board_id" string of the firmware. There are a couple of hard-coded board_id strings that are referenced. If neither of those match, NVRAM is queried to find out the running device's board_id. It's possible to verify the proper board ID is "U12H192T00_NETGEAR" by extracting the NVRAM parameters from a live device[3]. Even if we didn't have that information, we could still analyze a stock firmware, where we find the same string embedded in the header.

R6200-V1.0.0.28_1.0.24.chk

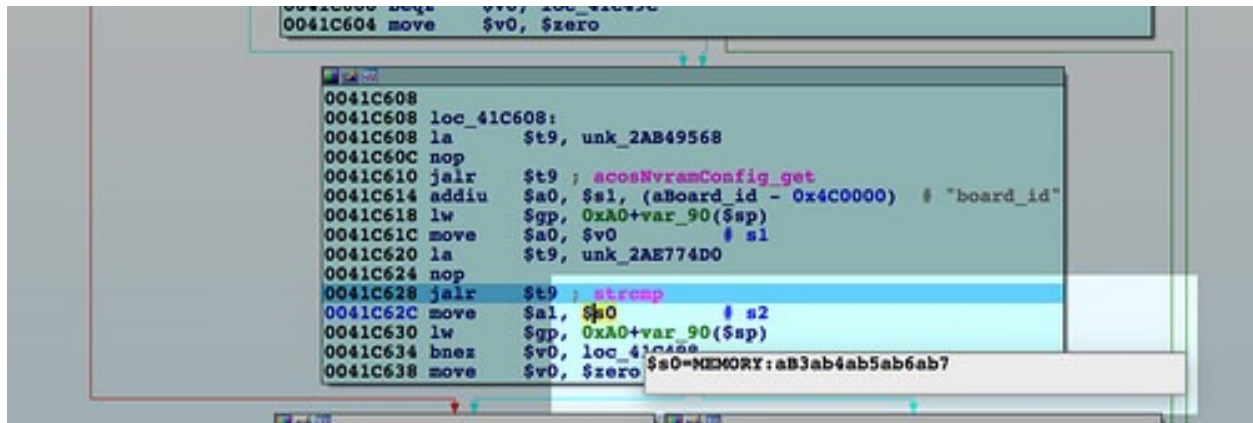
Save Copy Cut Paste Undo Redo Hex Hex search Go To Offset Find (Hex search)

| | | | | | | |
|--------|-------------|-------------|-------------|-------------|-------|--------------------|
| 000000 | 2A 23 24 5E | 00 00 00 3A | 01 01 00 00 | 1C 01 00 18 | F7 23 | *#\$^...:.....# |
| 000012 | 12 3C 00 00 | 00 00 00 87 | 10 00 00 00 | 00 00 F7 23 | 12 3C | .<.....#.< |
| 000024 | DE 65 09 5C | 55 31 32 48 | 31 39 32 54 | 30 30 5F 4E | 45 54 | .e.\U12H192T00_NET |
| 000036 | 47 45 41 52 | 48 44 52 30 | 00 10 87 00 | C0 39 E8 0E | 00 00 | GEARHDR0.....9.... |
| 000048 | 01 00 1C 00 | 00 00 04 45 | 14 00 00 00 | 00 00 5D 00 | 00 01 |E.....]... |
| 00005A | 00 86 D0 3B | 00 00 00 00 | 00 00 00 06 | A8 33 B8 F9 | A3 64 | ...;.3...d |
| 00006C | 4F 09 38 05 | 22 8E 52 10 | 58 DC 30 60 | E0 8C 0A 4A | 12 A5 | 0.8."R.X.0'...J.. |
| 00007E | E4 BC 8A 90 | 9C D0 49 97 | 51 93 8F E8 | 65 FA FA AF | 99 72 |I.Q...e....r |

| Type | Value |
|-----------------|--|
| 8 bit signed | 85 |
| 8 bit unsigned | 0x55 |
| 16 bit signed | 21809 |
| 16 bit unsigned | 0x5531 |
| 32 bit unsigned | 0x55313248 |
| 32 bit signed | 1429287496 |
| 64 bit unsigned | 0x5531324831393254 |
| 64 bit signed | 6138743052727562836 |
| BGR | 553132 |
| RGB | 553132 |
| binary | 01010101 00110001 00110010 01001000 00110001 00111001 00... |
| double (8 byte) | 2.40722E+102 |
| float (4 byte) | 1.21768E+13 |
| octal | 125 061 062 110 061 071 062 124 060 060 137 116 105 124 107 1... |
| string | U12H192T00_NETGEAR |

Hex Big Endian Overwrite Offset: 28 Selection: 12

As before, by looking at the pattern string that is compared, we can identify the offset into the header where the board_id should be placed.



```
$ ./buildfw.py find=b3Ab4Ab5Ab6Ab7Ab8A kernel.lzma
squashfs.bin
[@] Building firmware from input files: ['kernel.lzma',
'squashfs.bin']
[@] TRX crc32: 0x0ee839c0
[@] Creating ambit header.
[+] Building header without checksum.
[+] Calculating header checksum.
[@] Calculated header checksum: 0x840d0ddd
[+] Building header with checksum.
[@] Finding offset of b3Ab4Ab5Ab6Ab7Ab8A
[+] Offset: 40
```

The string b3Ab4Ab5Ab6Ab7Ab8A is located at offset 40.

It is worth noting that we suspected the header was variable length given the presence of a size field. The board_id is a string and is the last field in the header; it is likely responsible for the header's variable length.

At any rate, this is easy to add as a string section using Bowcaster. This is the last check in abCheckBoardID().

The Mystery Header So Far

Here's a diagram of what we know about the header so far.

| Byte | | |
|-------------|----------------|--------------------------------------|
| 0-3 | Magic: "*#\$^" | |
| 4-7 | Header Length | |
| 8-11 | | |
| 12-15 | | |
| 16-19 | | |
| 20-23 | | |
| 24-27 | | |
| 28-31 | | |
| 32-35 | | |
| 36-39 | | Header Checksum |
| 40-variable | | board_id "U12H192T00_NE TGEAR" |


```

SC.gadget_section(self.HEADER_SIZE_OFF,self.size,"Size field
representing length of header.")

SC.gadget_section(self.HEADER_CHECKSUM_OFF,checksum)
SC.string_section(self.BOARD_ID_OFF,self.board_id,
                  description="Board ID string.")
buf=OverflowBuffer(self.endianness,self.size,
                  overflow_sections=SC.section_list,
                  logger=logger)

def __checksum(self,header):
    data=str(header)
    size=len(data)
    chksum=LibAcosChecksum(data,size)
    return chksum.checksum

```

In the [next post](#) I'll discuss other functions that parse portions of the header. I'll show how to identify what fields get used where. By the end of the next installment we'll be able to generate a header sufficient to get our firmware image written to flash.

[1] Wah wah...Buffer overflow.

[2] I'm pretty sure it's [Fletcher32](#). I believe this because I asked [Dion Blazakis](#), and he thinks it is, and that dude is smart. Also I found a Fletcher32 [implementation](#) on Google Code by [Ange Albertini](#) that gives the same result as mine. And that guy is also smart.

[3] The NVRAM configuration can be extracted from /dev/mtd14. This, plus libnvram-faker is covered independently of this series, in [Patching, Emulating, and Debugging a Netgear Embedded Web Server](#)