



Flare-On 6: Challenge 6 – BMPHIDE.EXE

Challenge Author: Tyler Dean (@spresec)

Challenge Solution: Stephen Eckels (@ThorEckels)

This solution was written by Stephen Eckels after testing the challenge. He took two different approaches. Here is his story.

RECONNAISSANCE

This challenge starts with two files - `bmphide.exe` and `image.bmp`. The `image.bmp` file is a picture of mountains. From the looks of it, it appears to be a photo taken near Mt. Elbert, Colorado's tallest mountain.

Using a PE inspection tool such as CFF Explorer, we see that `bmphide.exe` is a .NET PE. My favorite tool for analyzing .NET samples is the open source tool dnSpy. Let's open the executable in dnSpy and see what we're dealing with. Looking at the executable, we see that all the class types have been renamed, they are simple single letter substitutions. Since we cannot gain much information from the names, we'll start by analyzing the entry point. We learn from the `Main` method that the application takes three arguments - all three are file paths. The first path is passed to the `Bitmap` object's constructor, so we know it should be an input image, and the second path is passed to `File.ReadAllBytes()`, so we know this is some sort of input binary as well. The final path is passed to `bitmap.Save()`, so we can assume this is an output path. In total we have an executable, that takes two inputs and one output. An example usage statement might like: `bmphide.exe ./inputImage.bmp ./inputText.bin ./outputImage.bmp`. From this, we assume this is some sort of steganography challenge that will combine the input text with the input image to produce some output image. We have our `image.bmp` file, so it's likely that we will extract the flag from this image. Let's continue our analysis.

OBFUSCATION PART 1

The first function called in `Main` is `Program.Init()`, so let's look there. We see that it loops through the `MethodInfo` structures of the 'A' class and calls `PrepareMethod()` on them, presumably to force JIT compilation on them as referenced here: [MSDN forcing JIT at runtime](#). After

the calls to `PrepareMethod()`, we see another loop. This time it loops through all methods in the `Program` class and gets their IL code as a byte array. Each byte array is passed to `D.a()`, and the result is compared with several constants. A byte array being converted to a constant for a comparison sounds like a hash, and if we visit `D.a()` we see that it does seem like a hashing method (Figure 1):

```
// Token: 0x00000002 RID: 2 RVA: 0x000020A0 File Offset: 0x000020A0
public uint a<T>(IEnumerable<T> byteStream)
{
    return ~byteStream.Aggregate(uint.MaxValue, (uint checksumRegister, T currentByte) => this.m_checksumTable[(int)((checksumRegister & 255u) ^ (uint)Convert.ToByte
(currentByte))] ^ checksumRegister >> 8);
}
```

Figure 1 – Notice the bitwise XOR operations and bitwise shifts in the `Aggregate`

We can then assume these hashes are being used to identify certain methods which are stored into local variables `m1-m4`. These locals then are passed to `A.VerifySignature()`, which we'll analyze next. In `VerifySignature()` the `PrepareMethod()` routine is again called on each method and we see pointer manipulation with the handles of the method. The last line dereferences the first method's pointer handle and assigns to it the second method's handle. This looks like it might not be an actual verify any signatures. Instead, it appears to be some sort of .NET internals method handle manipulation. Let's debug to check.

ANTI-DEBUG PART 1

Trying to debug to the method hash loop in `Init()`, we find that `dnSpy` encounters an unhandled exception (Figure 2):

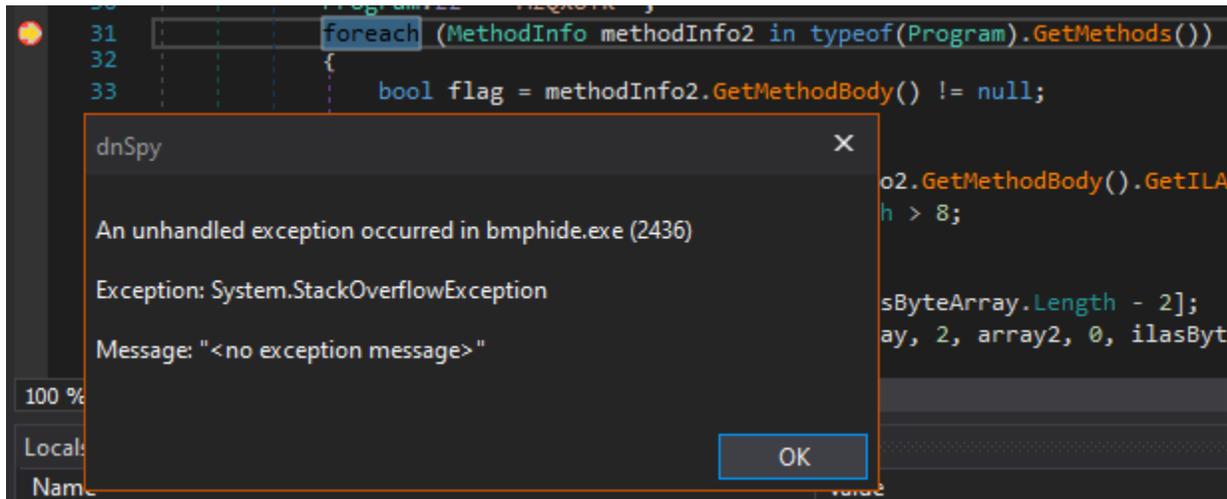


Figure 2 – Unexpected, Anti-Debug?

This is unfortunate, we don't yet know what causes this, and we can't debug to the interesting `VerifySignature()` method. Let's try to find the cause of this. We're already in `Init()`, the first method called, so we know something in here causes this exception, and we do see a

method called `CalculateStack()` before our interesting loop. Our exception message mentioned a stack overflow, so perhaps this is the cause. `CalculateStack()` calls `IdentifyLocals()` where we see a lot of interesting code. As seen in Figure 3, from a high level we see it calls `LoadLibrary`, `GetProcAddress`, `VirtualProtect`, `WriteIntPtr`, and manipulating function pointers; alarm bells start ringing, this is interesting.

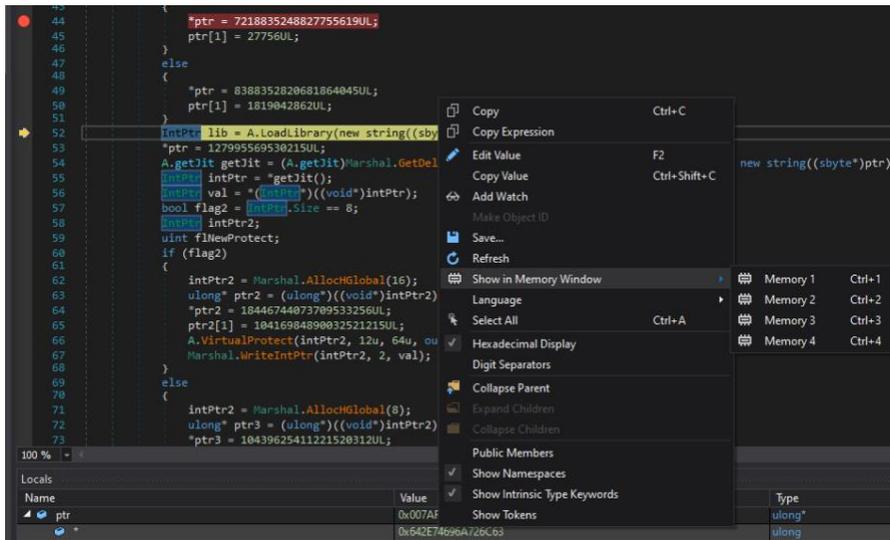


Figure 3 – Pointer Inspection

Analyzing further we see long constants being assigned into a `stackalloc` array. If we debug to this point, and inspect the `ptr` variable's value in memory by right clicking -> *Show in Memory Window*, we see the value points to the string `clrjit.dll`. We now know it's loading this DLL and resolving the `getJit` method (inspecting the pointer again shows the string `getJit`). If we keep following the logic, we land in another section where a pointer to an allocated area of memory is filled with a long constant, but this time it's not a string. We see a call to `VirtualProtect` to change the memory permissions to `PAGE_EXECUTE_READWRITE`, and we know that executable permission is weird on a data pointer. Maybe this value is code? If we convert the constant `10439625411221520312` to hex we get `90E0FFFFFFFFFFFFB8`, which when changing the endianness is `B8FFFFFFFFFFE090`. When disassembled, we see (Figure 4):

```

0:  b8 ff ff ff ff          mov     eax, 0xffffffff
5:  ff e0                   jmp     eax
7:  90                       nop

```

Figure 4 – Disassembly of pointer constant

The next line is a `WriteIntPtr` call with one byte offset into this assembly to overwrite the `0xFF`'s. The new value is a pointer to the `getJit` delegate, so we see this is a transfer stub back

to the original `getJit` method, aka a trampoline. Later down we see the `getJit` method overwritten by a Delegate to the `A.handler()`, and that the `originalDelegate` function is assigned to our trampoline assembly stub to perform the call back to the original. In summary, the code hooks the JIT method of the .NET runtime. Let's investigate this hook further and see if we can understand what it does and remove it.

JIT HOOK, OBFUSCATION PART 2, REMOVING ANTI-DEBUG

In the function `IncrementMaxStack()`, we see a method's `MetadataToken` is compared to two constants, `100663317` and `100663316` which in hex are `0x6000015` and `0x6000014` respectively. These are the `MetadataToken` values of the methods `Program.h()` and `Program.g()`. After checking each `MetadataToken` handler, the memory permissions are changed to allow writing. In both cases there are two values written to the IL code. We've identified another obfuscation technique; these methods are re-written at runtime to alter their logic. To see what these modifications do, we patch these bytes ourselves using `dnSpy`. This is accomplished by right clicking a method and selecting show instruction in hex editor. If we hover over the hex shown, we see an instruction array with an index that moves as we select different bytes. By navigating to offsets `0x17`, and `0x32` and modifying the byte values to `0x14`, we fixup method `Program.h()`. Performing the same manual patching on `Program.g()` we see the actual code that will execute at runtime. The patch with `0x14` being written in changed the logic in the `Program.h()` method to call `Program.g()` instead of `Program.f()`; `0x14` is for the `MetadataToken 0x6000014` of `Program.g()`. The next patch with the longer constants overwrote the constants in `Program.g()` to be `309030853` and `209897853`. Finalizing our analysis of the JIT hook we see that after these byte patches are performed, there is a call to `A.originalDelegate` which as we know points to our trampoline stub earlier to call the original JIT. To summarize, the JIT hook patches the contents of two methods. In one case, it changes the method called and the second case modifies constants. We now fully understand the JIT hook and can save our binary patches. Additionally, we remove the call to `CalculateStack()` by changing the call instruction to a NOP. To save the patches we click File->Save Module and go to MD Writer Options, and choose preserve All MD Tokens since we know this application is doing `MetadataToken` specific manipulation (we don't want to mess up anything we don't understand yet), and then clicking ok under Main. At this point the application is debuggable again with the hook removed.

METHOD OVERWRITE, OBFUSCATION PART 3

We assumed that the `VerifySignature` method was overwriting method handles. Now that the application is debuggable, let's validate this assumption. If we step into the method, we see that `m = Program.a()`, `m2=Program.b()`, `m3=Program.c()`, and `m4=Program.d()`. If we let the call finish, and then debug to the next point `Program.a()` or `Program.c()` is called we notice that the logic of `Program.b()` and `Program.d()` is where the code continues executing.

We conclude that the `VerifySignature` method is overwriting the method handles for the methods `Program.a()` and `Program.c()` to call `Program.b()` and `Program.d()`. Therefore, the methods `Program.a()` and `Program.c()` are never called and we can safely delete both these methods and patch any calls to the them with `Program.b()` and `Program.d()` instead. This concludes analysis of all methods in the `A` namespace. To summarize, we saw code hooking the .NET JIT function, so we assume this was a `HookManager` namespace or something of the sort. Since we've performed the equivalent patches the hook manager does at runtime, we can delete this entire namespace and re-save the module to focus on the actual logic in `Program`. For closure, the `D` namespace only had the 'a' method used to hash IL, so we can assume this was a `HashManager` of sorts and delete it as well.

SOLUTIONS

At this point we understand the protection mechanism of the application but not the logic. We'll now investigate a black-box solution that uses limited knowledge of logic to brute-force the decoding of the image at both stages. As well as a proper hand-crafted decoder that uses knowledge of the application's encoding routine.

BLACK BOX ANALYSIS

Looking at `Program.i()`, which is the last call before `bitmap.Save()` we can hope to gain a little understanding about the structure of the data this program writes out. We see a loop through the bounds of the image, getting each pixel per round. Each color element for each pixel is manipulated by mixing in data from the byte array read in. There are some obfuscated values though that make this more difficult to understand. Through quick debugging we see that `Program.j(27)` returns `0xF8`, and `Program.j(25)` returns `0xFC`. At this point we suspect a mask with the lower 3, 3, and 2 bits being used for data. Debugging the shift after the bitwise or operation we see that `Program.j(228)` return `0x07`, `Program.j(230)` returns `0x03`, and `Program.j(100)` returns `0x06`. This means our entire logic for encoding is a mask with data being encoded into the lower color bits where R is 3 bits, G is 3 bits, and B is 2 bits (Figure 5).

```
Color pixel = bm.GetPixel(i, j);
int red = ((int)pixel.R & 0xF8) | ((int)data[num] & 0x7);
int green = ((int)pixel.G & 0xF8) | (data[num] >> 0x03 & 0x07);
int blue = ((int)pixel.B & 0xFC) | (data[num] >> 0x06 & 0x03);
```

Figure 5 – Data Encoding

Let's attempt to identify a pattern in the output so that we can brute force. We know color and data are mixed, so if we zero out the color channel we are left only with data. To analyze the pattern of encoding we can input a black.bmp (all 0s) image the same size as the one given and feed it through the application with some known text as input to the byte array. To do this we fill an input text file with ASCII a characters, pass the black image plus this file to the application and generate output.bmp. We then write a small C# program to read this output.bmp file and mask off the color channels R,G,B so we can see the transformed data values. By cycling the first character of the input text from a-z we observe a repeating pattern of changes in the output.bmp.

```

a -> 02 03 01   m -> 02 03 02
b -> 02 05 01   n -> 02 05 02
c -> 02 07 01   o -> 02 07 02
d -> 02 01 00   p -> 03 01 01
e -> 02 03 00   q -> 03 03 01
f -> 02 05 00   r -> 03 05 01
g -> 02 07 00   s -> 03 07 01
h -> 02 01 03   t -> 03 01 00
i -> 02 03 03   u -> 03 03 00
j -> 02 05 03
k -> 02 07 03
l -> 02 01 02

```

Figure 6 – Pixel 0,0 for a-u

As seen in Figure 6, this output is pixel 0, 0 for multiple runs of the application with the first character in the input text file being the ASCII value on the left. Now notice that as we increase the ASCII value the value of the Green channel goes up, in a repeating pattern of 03, 05, 07, 01, repeat. At the repeat step of the Green, we notice the cycle of the Blue column iterates one, following a 01, 00, 03, 02, 01, repeat pattern. And then on the repeat step of Green the Red channel will increase its pattern by one step. This tells us that each character of input is essentially mapped to a unique 3 hex byte sequence, for the first character at least. We still need to investigate how other the second ASCII character maps and if it affects the encoding of

adjacent characters.

```
static void Main(string[] args)
{
    Bitmap bitmap = new Bitmap(args[0]);
    int pixelCount = bitmap.Width * bitmap.Height;

    int counter = 0;
    for (int j = 0; j < bitmap.Width; j++)
    {
        for (int i = 0; i < bitmap.Height; i++)
        {
            Color c = bitmap.GetPixel(j, i);
            int r = c.R & 7;
            int g = c.G & 7;
            int b = c.B & 3;
            if (counter++ > 20)
                break;

            Console.WriteLine(r.ToString("X2") + " " + g.ToString("X2") + " " + b.ToString("X2"));
        }
    }
}
```

Figure 5 – C# Output Pixel Data Decoder

To investigate the encoding of adjacent characters we change the input text file from a single letter a to two letters aa and run it through the application. In Figure 8 you can see that the encoding of the first character does not change, but the second character does not follow the same cycle we expected, if it did then the input ‘aa’ would have generated ‘02 03 01’ for both. Instead we see that the second characters do again follow a cycle, but their own cycle. We now understand that each position of ASCII input maps to a unique output, and that this holds for every position of the input text. We may have a duplicate in the pattern every now and then if we expand this the full non-ASCII byte range, but this is enough knowledge to brute force.

a -> 02 03 01	Before we write an application to brute force this, lets first establish the workflow by hand. To start, we will feed our black image and a text file with only a single character into the application and generate output a-Z (lower case and upper). We will then run our C# data decoder on the given challenge image, <code>image.bmp</code> . This will give us just the lower data bit mappings for whatever the original data was that generated the image. From here we will attempt to pair this data decoder output, to the mappings we generated from the black image run. If we find a match between the black image run, and challenge image, then we can see which ASCII value generated that mapping and recover the input character by character.
a -> 01 00 00	
a -> 02 03 01	
b -> 01 06 00	
a -> 02 03 01	
c -> 01 04 00	
a -> 02 03 01	
d -> 01 02 01	
a -> 02 03 01	
e -> 01 00 01	
a -> 02 03 01	
f -> 01 06 01	

Figure 6 – Adjacent Character Runs

Let's walk through the first three values we recover with this approach. We recover 'B' in first position (00 05 01), 'M' in second position (03 00 03), and '6' in third position ('06 05 02'). This is the magic number for a BMP, telling us that inside this bmp another bmp has been mixed in, and that we need to expand our solution the entire byte range since bmp is a binary and not ASCII text.

We now also know that our manual workflow should be to match outputs from the black image run, to the output of the data decoder run on the challenge image. So, we will write a brute force application that will generate all possible outputs for each character position, and match pixel by pixel to the output of the challenge image. The challenge image is 1664x1248 = 2,076,672 pixels. Our embedded image is somewhere around that size, and we need to generate 0-255 inputs for each of those pixels to attempt a match. The only scalable way to do this is to create 255 binary files, filled with 0x00-0xFF, each being a solid fill of only that value. Each of these binary files will then be sent through the application with the black image, giving an output image. We keep track of the input -> output image pairs, when a particular data value in a pixel from the challenge image matches one of the values at the same pixel in an output image, we can see which input file created that output image and hence the original binary value that belongs at that pixel position. So in summary, for each data value from pixel I,J in challenge image search all of the 255 output files at pixel I,J, if there is a match in one of the output files, lookup the input file that generated it to recover the binary data, replace pixel I,J of challenge image with this value. A screenshot of the brute force tool in action is shown in Figure 9.

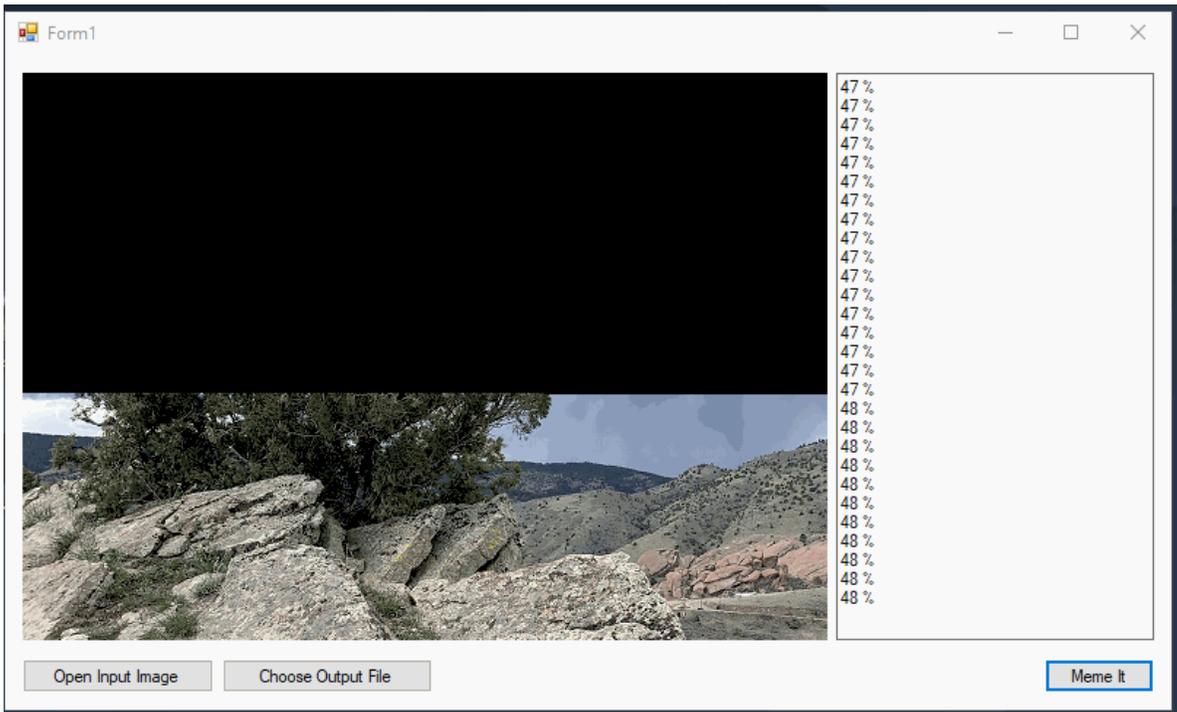


Figure 7 – Brute Forcing Stage 1 Animation.

The code I used to brute force is shown in Figure 10.

```

1. Bitmap bitmap = new Bitmap(inputImage);
2. int pixelCount = bitmap.Width * bitmap.Height;
3.
4. for (int i = 0; i < 256; i++)
5. {
6.     if (System.IO.File.Exists(cwd + "output" + i + ".bmp"))
7.         continue;
8.
9.     byte[] tmp = Enumerable.Repeat((byte)i, pixelCount).ToArray();
10.    System.IO.File.WriteAllBytes(cwd + "input" + i, tmp);
11.    var proc = new Process
12.    {
13.        StartInfo = new ProcessStartInfo
14.        {
15.            FileName = cwd + "bmphide.exe",
16.            Arguments = cwd + "black.bmp " + cwd + "input" + i + " " + cwd +
"output" + i + ".bmp",
17.            UseShellExecute = false,
18.            RedirectStandardOutput = true,
19.            CreateNoWindow = true
20.        }
21.    };
22.    proc.Start();
23.    Console.WriteLine("Waiting for output" + i);
24.
25.    if (i % 10 == 0 || i >= 244)
26.        proc.WaitForExit();
27. }
28.
29. Console.WriteLine("Making haystacks");
30. Bitmap[] imgArr = new Bitmap[256];
31. for (int i = 0; i < 256; i++)
32. {
33.     Console.WriteLine("Loading haystack " + i);
34.     imgArr[i] = new Bitmap(cwd + "output" + i + ".bmp");
35. }
36. Console.WriteLine("Starting Search");
37. byte[] output = new byte[pixelCount];
38. int outIdx = 0;
39.
40. for (int j = 0; j < bitmap.Width; j++)
41. {
42.     for (int i = 0; i < bitmap.Height; i++)
43.     {
44.         Color c = bitmap.GetPixel(j, i);
45.         int r = c.R & 7;
46.         int g = c.G & 7;
47.         int b = c.B & 3;
48.
49.         for (int k = 0; k < 256; k++)
50.         {
51.             Bitmap img = imgArr[k];
52.             Color haystack = img.GetPixel(j, i);
53.             if (r == haystack.R && g == haystack.G && b == haystack.B)
54.             {
55.                 output[outIdx++] = (byte)k;
56.                 break;
57.             }
58.         }
59.     }
60. }
61.
62. System.IO.File.WriteAllBytes(outputImage, output);

```

Figure 10 – Brute force code

This will decode the first embedded image correctly. The output of this first extraction will need to feed through the brute forcer again, and will produce a slightly corrupted image, likely due to the size of the second image being different from the original image (Figure 11).

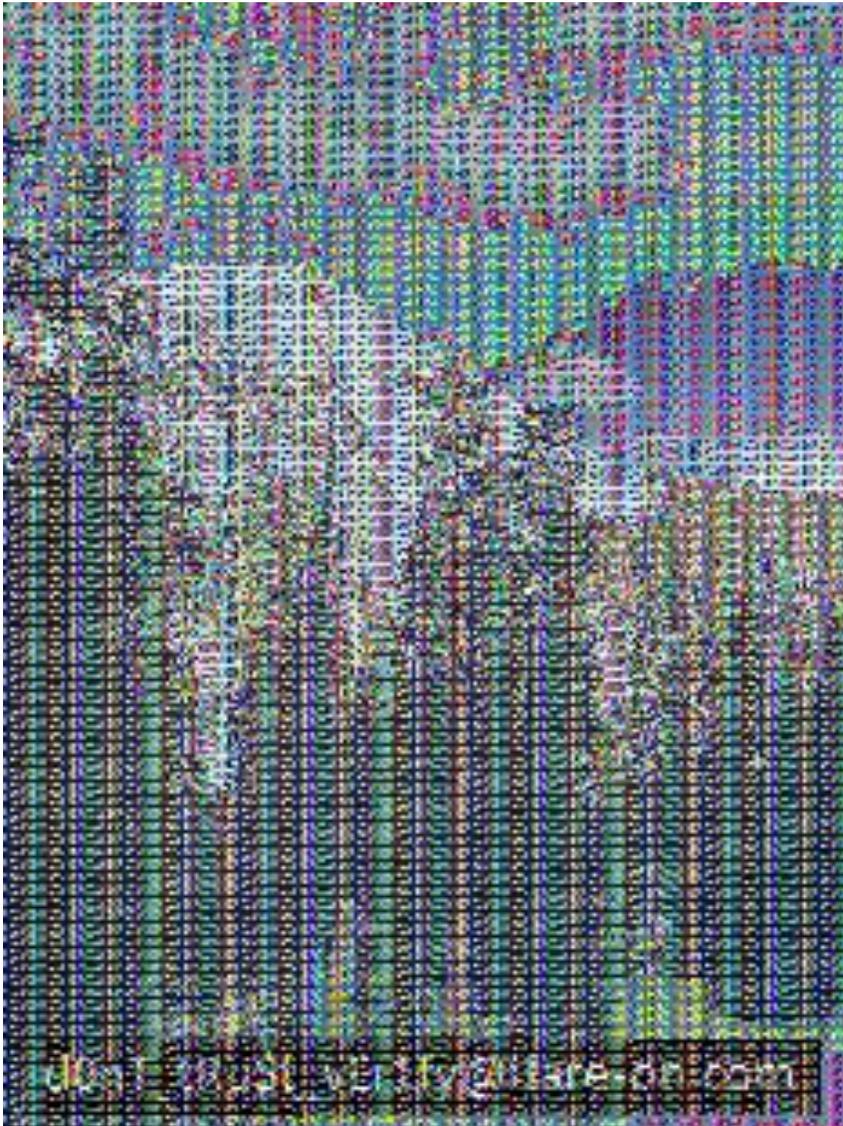


Figure 11 – Brute force result

But if you look closely, at the bottom we can almost see the flag. To fix we just multiply every pixel value by 4, which helps smooth out the noise (Figure 12 and Figure 13).

```

1. #include <fstream>
2. int main() {
3.     std::fstream file = std::fstream("C:\\outputhack.bmp", std::ios::out |
std::ios::binary);
4.     // don't start at beginning, to skip BMP header (overestimate)
5.     for (unsigned long long i = 3*20 + 2; i < stage2_len ; i++) {
6.         stage2[i] *= 4;
7.     }
8.     file.write((char*)stage2, stage2_len);
9.     file.close();
10. }

```

Figure 12 – Fix-up bitmap

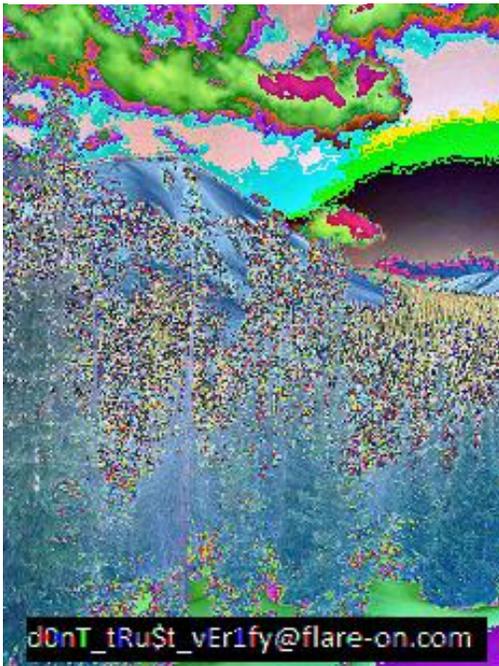


Figure 8 – Smoothed Output

ENCODER ROUTINE ANALYSIS & DECODER SOLUTION

After going through the brute force solution, let's walk through the way the challenge was originally intended to be solved. We already understand how the file is written out, with lower bits being used for data, but we do not yet understand the transformation applied to the input binary array. The previous solution simply brute forced the input, output pairs in order to

transformation extract the output. Let's take a look at this transformation step in a bit more detail. The function of interest is `Program.h()`, which loops over all data and does a series of calls manipulating the array byte by byte. We start our analysis with `Program.g()`. We see it takes as input an index, and inside multiplies that by two constants. Next it calls the `Program.e()` method and returns the result. This appears to be a pseudo random number generator of sorts.

Let's look closer at the `Program.e()` method. This method iterates over each bit of both input byte values. If both bits are equal, a bitwise NOT is applied to the bit and a bitwise AND is applied to the output. If they are not equal, a bitwise OR is applied. We see that the cases where the bits are not equal, the result is always 1. If they are equal, the result is always 0. This sounds a lot like XOR. The method `Program.g()` is a pseudo random number generator and `Program.e()` is XOR.

We see the XOR used again in `Program.h()`, so we continue to the next unknown call to `Program.b()`. This method is called with the result of the XOR and the number 7 as arguments. There's a loop that execute 0-7 and at each round it masks with 128 (b1000 0000) to get the most significant bit, and then divides by 128 to shift that MSB to the least significant bit position, storing the result in `b2`. The input is then multiplied by 2, which is the same as a shift left by one position, masked with an and to keep it in byte range, and then adding `b2` back on. In English *Get the MSB and store it, shift everything else one position left, and place the stored value back onto the LSB, repeat up to 7 times*, aka a rotate left by 7.

The next unknown function is `Program.d()`, immediately we see it's similar to the ROL we just reversed, but instead it took 3 as it's argument and the operations are different. It masks the first bit with 1 to get the LSB, then multiplies by 128 to left shift it in the MSB position, storing in `b2`. The input is then divided by 2 to right shift, masked to keep in byte range, then `b2` is added back on. In English *Get the LSB and store it, shift everything else one position right, and place the stored value back onto the MSB, repeat up to 3 times*, aka a rotate right by 3.

```

1. public static byte[] h(byte[] data)
2. {
3.     byte[] array = new byte[data.Length];
4.     int num = 0;
5.     for (int i = 0; i < data.Length; i++)
6.     {
7.         int num2 = (int)Program.Hash(num++);
8.         int num3 = (int)data[i];
9.         num3 = (int)Program.Xor((byte)num3, (byte)num2);
10.        num3 = (int)Program.ROL((byte)num3, 7);
11.        int num4 = (int)Program.Hash(num++);
12.        num3 = (int)Program.Xor((byte)num3, (byte)num4);
13.        num3 = (int)Program.ROR((byte)num3, 3);
14.        array[i] = (byte)num3;
15.    }
16.    return array;
17. }

```

Figure 14 – Program.h with readable method names

To decode, we need to undo the rotates and then XOR again. We can recover the random number values by re-using the same routine and feeding the indices through it in the same order. Since we're doing the decoding in reverse order, we need to take care to call `Program.Hash()` (our label for the random generator) in the same order as original but use the values in reverse order. We plan to ROL by 3, then XOR with the output of the second call to `Program.Hash()`, ROR 7, then XOR with the output of the first `Program.Hash()`, and the resultant value should be the original binary value encoded in the image at that pixel. We need to take care to only run this decode routine on the data bits of the image, so we also must mask the R,G,B channels of each pixel by 7, 7, and 3 and re-assemble the channels into a byte. The full code is shown in Figure 15.

```

1. class Program
2. {
3.     public static byte hashidx(int idx)
4.     {
5.         byte b = (byte)((long)(idx + 1) * (long)((ulong)309030853));
6.         byte k = (byte)((idx + 2) * 209897853);
7.         return (byte)(b ^ k);
8.     }
9.
10.    public static byte Rol(byte b, int r)
11.    {
12.        for (int i = 0; i < r; i++)
13.        {
14.            byte b2 = (byte)((b & 128) / 128);
15.            b = (byte)((b * 2 & byte.MaxValue) + b2);
16.        }
17.        return b;
18.    }
19.
20.    public static byte Ror(byte b, int r)
21.    {
22.        for (int i = 0; i < r; i++)
23.        {
24.            byte b2 = (byte)((b & 1) * 128);
25.            b = (byte)((b / 2 & byte.MaxValue) + b2);
26.        }
27.        return b;
28.    }
29.
30.    static void Main(string[] args)
31.    {
32.        Bitmap bitmap = new Bitmap(args[0]);
33.        int pixelCount = bitmap.Width * bitmap.Height;
34.
35.        Console.WriteLine("Starting Search");
36.        byte[] output = new byte[pixelCount];
37.        int outIdx = 0;
38.
39.        int hashCtr = 0;
40.        for (int j = 0; j < bitmap.Width; j++)
41.        {
42.            for (int i = 0; i < bitmap.Height; i++)
43.            {
44.                Color c = bitmap.GetPixel(j, i);
45.                // Mask off image data
46.                int r = c.R & 7;
47.                int g = c.G & 7;
48.                int b = c.B & 3;
49.
50.                // Re-assemble encoded char to byte
51.                byte orig = (byte)b;
52.                orig <<= 3;
53.                orig |= (byte)g;
54.                orig <<= 3;
55.                orig |= (byte)r;
56.
57.                // Get the two hashed idxs
58.                byte g1 = hashidx(hashCtr++);
59.                byte g2 = hashidx(hashCtr++);
60.
61.                byte newByte = Rol(orig, 3);
62.                newByte = (byte)(newByte ^ g2);
63.                newByte = Ror(newByte, 7);
64.                newByte = (byte)(newByte ^ g1);
65.
66.                output[outIdx++] = newByte;
67.            }
68.        }
69.
70.        System.IO.File.WriteAllBytes(args[1], output);
71.    }
72. }

```

Figure 15 – Final recovery tool

The final resulting image shown in Figure 16 contains the key (“d0nT_tRu\$t_vEr1fy@flare-on.com”):

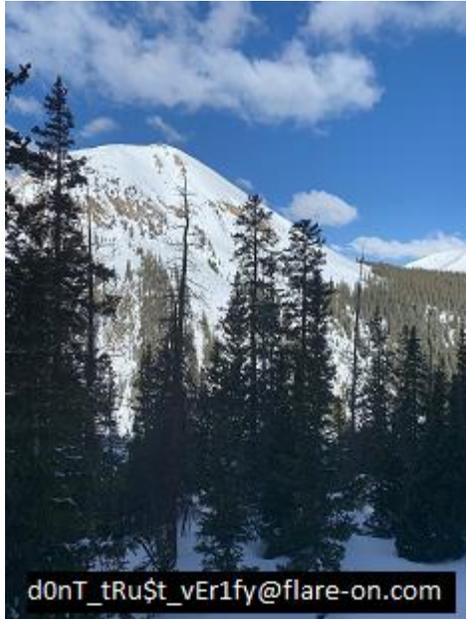


Figure 16 – Challenge key