New issue

Heap buffer over-read in avi_parse_input_file() when parsing truncated hdrl strh chunk #24

Open

BreakingBad6 opened 2 weeks ago · edited by BreakingBad6 Edits

Minimal harness

```
#include <stdio.h>
#include "../example/theora/win32/experimental/transcoder/avi2vp3/avilib.h"

int main(int argc, char **argv) {
    const char *path = "poc/avilib_hdrl_oob/avilib_hdrl_short_strh.avi";
    avi_t *avi;

    if (argc > 1) {
        path = argv[1];
    }

    printf("opening: %s\n", path);
    avi = AVI_open_input_file((char *)path, 1);
    if (!avi) {
        AVI_print_error("AVI_open_input_file failed");
        return 1;
    }

    printf("opened successfully\n");
    AVI_close(avi);
    return 0;
}
```

Build
clang -g -O0 -fsanitize=address -fno-omit-frame-pointer -
Iexample/theora/win32/experimental/transcoder/avi2vp3 poc/avilib_hdrl_asan_harness.c
example/theora/win32/experimental/transcoder/avi2vp3/avilib.c -o avilib_hdrl_asan_harness

Run
ASAN_OPTIONS=detect_leaks=0:abort_on_error=1:symbolize=1 ./avilib_hdrl_asan_harness poc/avilib_hdrl_oob/avilib_hdrl_short_strh.avi

ASan result
The crash is an AddressSanitizer heap-buffer-overflow read in:

memcpy(AVI->compressor,hdrl_data+i+4,4)
The allocated heap region is created at:

hdrl_data = (unsigned char *) malloc(n);
I will attach the ASan output screenshot separately.

Impact
This issue allows a malformed AVI file to trigger a heap out-of-bounds read and crash the parser. Based on current analysis, this appears to be a denial-of-service / parser robustness issue caused by missing bounds checks in hdrl sub-chunk parsing.

Suggested fix
Before reading any nested hdrl chunk fields:

Verify that at least 8 bytes remain for the chunk header.
Verify that the declared chunk length fits within the remaining hdrl_data buffer.
For strh / strf, verify that the remaining bytes are sufficient for all fixed-offset accesses before reading fields such as +4, +14, +16, +20, +24, or +32.
Abort parsing on malformed or truncated sub-chunks.

```
    #6 0x560e56d68354 in _start (/mnt/c/Users/18320/Desktop/新农计划/avilib_hdrl_asan_harness+0x1f354) (BuildId: 7d28785
3e9f01b2dc811b9bee12fbdec4901f8ea)

0x60200000001c is located 0 bytes to the right of 12-byte region [0x602000000010,0x60200000001c)
allocated by thread T0 here:
    #0 0x560e56deb19e in __interceptor_malloc (/mnt/c/Users/18320/Desktop/新农计划/avilib_hdrl_asan_harness+0xa219e) (Bu
ildId: 7d287853e9f01b2dc811b9bee12fbdec4901f8ea)
    #1 0x560e56e2dc65 in avi_parse_input_file /mnt/c/Users/18320/Desktop/新农计划/example/theora/win32/experimental/tran
scoder/avi2vp3/avilib.c:1169:43
    #2 0x560e56e2d763 in AVI_open_input_file /mnt/c/Users/18320/Desktop/新农计划/example/theora/win32/experimental/trans
coder/avi2vp3/avilib.c:1093:3
    #3 0x560e56e25f6b in main /mnt/c/Users/18320/Desktop/新农计划/poc/avilib_hdrl_asan_harness.c:13:11
    #4 0x762bb3029d8f in __libc_start_call_main csu/../sysdeps/nptl/libc_start_call_main.h:58:16

SUMMARY: AddressSanitizer: heap-buffer-overflow (/mnt/c/Users/18320/Desktop/新农计划/avilib_hdrl_asan_harness+0xa14c6) (
BuildId: 7d287853e9f01b2dc811b9bee12fbdec4901f8ea) in __asan_memcpy
Shadow bytes around the buggy address:
  0x0c047fff7fb0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  0x0c047fff7fc0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  0x0c047fff7fd0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  0x0c047fff7fe0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  0x0c047fff7ff0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
=>0x0c047fff8000: fa fa 00[04]fa fa fa fa fa fa fa fa fa fa fa fa
  0x0c047fff8010: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
  0x0c047fff8020: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
```

```
  0x0c047fff7fe0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  0x0c047fff7ff0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
=>0x0c047fff8000: fa fa 00[04]fa fa fa fa fa fa fa fa fa fa fa fa
  0x0c047fff8010: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
  0x0c047fff8020: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
  0x0c047fff8030: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
  0x0c047fff8040: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
  0x0c047fff8050: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
Shadow byte legend (one shadow byte represents 8 application bytes):
  Addressable:           00
  Partially addressable: 01 02 03 04 05 06 07
  Heap left redzone:       fa
  Freed heap region:       fd
  Stack left redzone:      f1
  Stack mid redzone:       f2
  Stack right redzone:     f3
  Stack after return:      f5
  Stack use after scope:   f8
  Global redzone:          f9
  Global init order:       f6
  Poisoned by user:        f7
  Container overflow:      fc
  Array cookie:            ac
  Intra object redzone:    bb
  ASan internal:           fe
  Left alloca redzone:     ca
  Right alloca redzone:    cb
==1223==ABORTING
Aborted (core dumped)
```

BreakingBad6 mentioned this 2 weeks ago

Stack buffer overflow in txtin_process_texml (load_text.c:3844) gpac/gpac#3467

BreakingBad6 2 weeks ago Author • • •

the poc:

新建文件夹 (10).zip

Metadata

Assignees

No one assigned

Labels

No labels

Type

No type

Projects

No projects

Milestone

No milestone

Relationships

None yet

Development

Code with agent mode

No branches or pull requests

Participants