# riscure

driving your security forward

# **black hat** EUROPE 2021

november 10-11, 2021

BRIEFINGS

Achieving Linux Kernel Code Execution Through A Malicious USB Device

Martijn Bogaard Principal Security Analyst <u>martijn@riscure.com</u> @jmartijnb Dana Geist Senior Security Analyst Əgeistdana



#### Our agenda

- Why USB based attacks?
- The forgotten vulnerability (a.k.a CVE-2016-2384)
- Exploitation approach
- Demo
- Discussion / Final thoughts

# Why USB based attacks?





#### Why USB based attacks?



## **The forgotten vulnerability** CVE-2016-2384





#### CVE-2016-2384: what is it about?

The forgotten vulnerability

- Double free in the Linux kernel found<sup>1</sup> by Andrey Konovalov
- Targets the USB MIDI subsystem
- PoCs demonstrating:
  - Denial of Service
  - Code execution in the kernel turning the double-free into a use-after-free by:
    - Unprivileged code execution (i.e.: syscall interface)
    - Sockets (i.e.: allocation of SKBs)
- Affected major Linux based systems and distributions such as Ubuntu Fedora and CentOS assuming physical access



#### CVE-2016-2384: why do we still care?

The forgotten vulnerability



#### DeviceX characteristics

The forgotten vulnerability

- Widely used device
- Very minimal Linux-based system
  - Kernel early 4.x
- Actively backporting security fixes
- Locked down (attack surface reduction)
  - No crash logs, no serial output

The actual device is not relevant for this talk ;-)

USB (not used)	وممر	Network (1 service)
	Sandboyed	
	browser	



# CVE-2016-2384: how is the double-free triggered?

• A USB MIDI device is connected to the target device

• The device's configuration is not standard





#### USB probing process

USB crash course

During probing all communication is initiated by the host:





## CVE-2016-2384: how is the double-free triggered?

The forgotten vulnerability

- Since the device is of type MIDI, a specific function is called: \_\_\_\_snd\_usbmidi\_create()
- The function allocates a structure on the heap





## CVE-2016-2384: how is the double-free triggered?

The forgotten vulnerability

- A (deliberate) error in the MIDI configuration is encountered:
  - Causing the 1<sup>st</sup> free: **snd\_usbmidi\_free(umidi)**
  - But also returns an error which makes the entire probing process fail
- As part of the cleanup process a free is executed on the same memory location
  - Causing the 2<sup>nd</sup> free: snd\_usbmidi\_free(umidi)



#### To SLAB or to SLUB

SLUB allocator

- SLUB is the default allocator in Linux (since 2.6.23)
  - Implements the nitty gritty details of the kernel allocations and deallocations
- SLUB groups allocated chunks into *slabs* 
  - A slab is a collection of objects of the same (rounded) size
  - The freelist is arranged as a simple linked list (next ptr = only metadata)
  - When allocating a new chunk, the first object in the list will be removed from the list and its pointer returned

This specific behavior shapes the freelist in a very specific way after a double-free occurs!



#### CVE-2016-2384: what happens with the heap?

SLUB allocator





#### CVE-2016-2384: the fix

The forgotten vulnerability

# **Our exploitation approach**





#### USB probing process

USB crash course

- Devices are constrained to reply to the host requests and cannot:
  - Initiate arbitrary communication
  - Send arbitrary data in USB packets

#### $\rightarrow$ Limited exploitation primitives!

#### The midi object

Exploitation approach

- Let's have a look at our MIDI object
- Struct size is between 256 and 512
- This means we need to focus on the slab-512
  - Big enough to hold a reasonable payload
  - In general this is a low activity slab

This increases our chances of winning the race!







Exploitation approach

- So how can turn a double free into something useful?
  - Remember, we have a loop in our freelist
- All allocations on the same CPU + Slab get the same chunk

ightarrow Often results in a kernel panic within seconds



But every allocation updates the freelist ptr with the content of the first 8 bytes of the chunk!



Exploitation approach

• All allocations on the same CPU + Slab get the same chunk



But every allocation updates the freelist ptr with the content of the first 8 bytes of the chunk!

 $\rightarrow$  2 consecutive allocations with control over the content will result in freelist ptr control



Exploitation approach

• All allocations on the same CPU + Slab get the same chunk



But every allocation updates the freelist ptr with the content of the first 8 bytes of the chunk!

- → 2 consecutive allocations with control over the content will result in freelist ptr control
- → The 3rd allocation gives then an arbitrary write primitive!



Exploitation approach

- 512-slab chunk
- We need 3 kmallocs in a row
- Controlled data in the first and last allocation
  - 1<sup>st</sup> chunk: contains pointer to **arbitrary\_mem\_location**
  - 3<sup>rd</sup> chunk: payload written to **arbitrary\_mem\_location**

So where to find such primitive...?

• Let's dive in another USB driver

#### **USB HID**

Exploitation approach

- Human Interface Device protocol
  - A generic protocol for keyboards, mouse's, game controllers, etc.
  - Describes a device as a series of inputs and outputs
- Uses HID and HID REPORT descriptors to describe the device functions
  - Report descriptor is a variable length blob of data (up to 4 KiB)

0x05, 0x01, // USAGE PAGE (Generic Desktop) // USAGE (Mouse) // COLLECTION (Application) USAGE (Pointer) 11 COLLECTION (Physical) 11 USAGE PAGE (Button) 11 USAGE MINIMUM (Button 1) 11 USAGE MAXIMUM (Button 3) 11 LOGICAL\_MINIMUM (0) 11 LOGICAL\_MAXIMUM (1) 0x25, 0x01, 11 REPORT COUNT (3) 11 REPORT SIZE (1) 0x75, 0x01, 0x81, 0x02, 11 INPUT (Data, Var, Abs) 0x95, 0x01, 11 REPORT COUNT (1) 0x75, 0x05, 11 REPORT\_SIZE (5) 11 INPUT (Cnst,Var,Abs) 11 USAGE PAGE (Generic Desktop) 0x09, 0x30, 11 USAGE (X) 0x09, 0x31, 11 USAGE (Y) 0x15, 0x81, 11 LOGICAL MINIMUM (-127) 0x25, 0x7f, 11 LOGICAL MAXIMUM (127) 11 REPORT SIZE (8) 0x95, 0x02, 11 REPORT COUNT (2) 0x81, 0x06, 11 INPUT (Data, Var, Rel) 11 END COLLECTION // END COLLECTION

0x09, 0x02,

0xa1, 0x01,

0x09, 0x01,

0xa1, 0x00,

0x05, 0x09,

0x19, 0x01,

0x29, 0x03,

0x15, 0x00,

0x95, 0x03,

0x81, 0x03, 0x05, 0x01,

0x75, 0x08,

0xc0,

0xc0



#### USB HID - Probing

```
static int usbhid_parse(struct hid_device *hid) {
        [...]
        if (!rsize || rsize > HID_MAX_DESCRIPTOR_SIZE) {
                dbg_hid("weird size of report descriptor (%u)\n", rsize);
                return - EINVAL;
        ret hid_get_class_descriptor(dev, interface->desc.blnterfaceNumber, HID_DT_REPORT, rdesc, rsize);
        if (ret < 0) { .. }
```



#### USB HID - Probing

```
int hid_open_report(struct hid_device *device) {
    [..]
```

We get the report descriptor and make 2 copies of it! → Arbitrary write

```
start = device->dev_rdesc;
size = device->dev_rsize;
```

```
buf = kmemduplstart, size, GFP_KERNEL);
[..]
start = kmemduplstart, size, GFP_KERNEL);
```

```
kfree(buf);
```

```
device->rdesc = start;
device->rsize = size;
```

```
parser = vzalloc(sizeof(struct hid_parser));
```



# **Exploitation in the dark**





#### Development environment

Exploitation Approach

Even if the bug is there...

- How do we develop an exploit without any means to debug?
  - No crash logs, no serial port, nothing
- How do we make sure we can predictably win the race?
- What is our payload supposed to look like?

Aaand the winner is:





#### Development environment: QEMU

- Spend a lot of time making an environment close to the real device
  - Benefit: Full control & ability to debug the attack
    - Critical to build a deep understanding of all the steps
- Challenges:
  - An obvious consequence is that the target binary will be different
    - At least ensure all critical code paths & data structures are identical!
  - Also the device's activity level may vary
    - More activity == more chance of losing the race
  - How accurate is the emulation?



#### Development environment: QEMU

Exploitation Approach

- How to test the attack?
  - We can use the QEMU emulated devices
    - usb-mouse  $\rightarrow$  MIDI device (with invalid configuration)
    - usb-tablet  $\rightarrow$  Delivers payload through HID report descriptor

Major surprise:

When we did it on the real device, it almost immediately worked!



#### Payload delivery method

- No off-the-shelf device will:
  - Cause the double-free using a MIDI device
  - Allow us to deliver the payload in an HID report descriptor





#### Payload delivery method

- No off-the-shelf device will:
  - Cause the double-free using a MIDI device
  - Allow us to deliver the payload in an HID report descriptor









#### Where to hijack the code?

Exploitation Approach

.text:FFFF00000886F1B0	device = X22	; hid_device *
.text:FFFF00000886F1B0	BL	_mcount
.text:FFFF00000886F1B4	LDR	W0, [device,#0x1BC0]
.text:FFFF00000886F1B8	TBNZ	W0, #1, loc_FFFF00000886F3C0
<pre>.text:FFFF00000886F1BC</pre>	LDR	X0, [device] ; src
.text:FFFF00000886F1C0	CBZ	X0, loc_FFFF00000886F3E8
.text:FFFF00000886F1C4	LDR	W1, [device,#8]
.text:FFFF00000886F1C8	MOV	W2, #0x24000C0 ; gfp
.text:FFFF00000886F1D0	STR	W1, [X29,#0x70+len]
.text:FFFF00000886F1D4	MOV	W1,W1 ; len
.text:FFFF00000886F1D8	BL	kmemdup
.text:FFFF00000886F1DC	MOV	X20, X0
.text:FFFF00000886F1E0	CBZ	X0, loc_FFFF00000886F428
.text:FFFF00000886F1E4	LDR	X0, [device,#0x1BB0]
.text:FFFF00000886F1E8	LDR	X3, [X0,#0x60]
.text:FFFF00000886F1EC	CBZ	X3, loc_FFFF00000886F308
.text:FFFF00000886F1F0	ADD	X2, X29, #0x5C ; '\'
.text:FFFF00000886F1F4	MOV	X1, X20
.text:FFFF00000886F1F8	MOV	X0, device
.text:FFFF00000886F1FC	BLR	X3
.text:FFFF00000886F200		
.text:FFFF00000886F200	loc_FFFF00000886F200	; CODE XREF: hid_open_report+17C↓j
.text:FFFF00000886F200	LDR	W1, [X29,#0x70+len] ; len
.text:FFFF00000886F204	MOV	W2, #0x24000C0 ; gfp
.text:FFFF00000886F20C	BL	kmemdup
.text:FFFF00000886F210	MOV	X19, X0
.text:FFFF00000886F214	MOV	X0, X20 ; x
.text:FFFF00000886F218	BL	kfree
.text:FFFF00000886F21C	CBZ	X19, loc_FFFF00000886F428
.text:FFFF00000886F220	LDR	W1, [X29,#0x70+len]
.text:FFFF00000886F224	MOV	X0, #0x18100 ; size
.text:FFFF00000886F22C	STR	X19, [device,#0x10]

33



#### Payload design

- Started with testing different payloads in QEMU
  - First payload only called called printk() & crashed
  - Next version called usb\_get\_string() (observable from the outside)
  - Final version could run an indefinite amount of times!















#### Run shell commands

- Standard method to run arbitrary commands is using run\_cmd()
  - However: cannot run from interrupt context
    - usb\_hub\_wq runs under this context
- Alternative:
  - Use **system\_wq** to schedule start of new process
- orderly\_reboot() / orderly\_shutdown()
  - Both have writeable commands they are going to execute ightarrow overwrite & call



#### Run shell commands

- What shall we run?
  - Existing binary in rootfs (e.g. wget, nc, ...)
  - When a USB stick is auto-mounted ightarrow Run binary from USB stick (e.g. reverse shell)
  - Otherwise?
    - Make device node for USB stick
    - Mount USB stick
    - Then run the binary
- What about SELinux preventing to spawn processes?
  - Disable it! (selinux\_enforcing = 0, patching the policydb, unhooking the LSM hooks)

# **Demo time!**



Listening on [0.0.0.0] (family 2, port 1337)



# **Discussion**







#### Attack challenges

- Winning the race
  - Midi and HID probe need to happen before kernel panics due to corruption
  - Midi and HID probe might happen on different cores due to scheduling
    - High chance (~50-80%) of winning the race when system is idle
- Cache behavior
- Exploit mitigations



#### Exploit mitigations

- ASLR?
  - Leak kernel pointer (e.g. uninitialized memory vulnerability)
  - Alternatively brute-force destination address (only when crashing the device several times is acceptable)
- Write protected kernel code / rodata
  - Data-only attack / targeting code of priv. umode process through physmap
  - Might be hard to not crash the kernel due to heap corruption!
  - Also: demo device discards writes but caches them until write-back
    - Still exploitable!
- Freelist hardening / randomization
  - Might be truly making this attack hard / infeasible



#### Applicability

- What about similar bugs?
  - Likely exploitable using this approach when it results in a double free on a low-activity slab
- Exploitation steps may differ
  - Architecture specific characteristics (cache, etc.)
  - Different Linux kernel configurations
  - Implemented exploit mitigations
  - Vendor specific customizations



#### Exploitation Requirements

[ ] Linux based device vulnerable to CVE-2016-2384[ ] Physical access

[-] Unprivileged code execution[-] Networking interface

[+] Low activity slab



#### Takeaways

- Fixed vulnerabilities upstream can take a long time to propagate
- Invest in building debug capabilities for your target!
  - Either through emulation or other means
- USB attacks are powerful & physical access might be all you need
  - In some cases it's your only attack vector!