



A Disaster Caused By A Bug: A Black Box Escape Of Qemu Based On The USB Device

flyyy@Chaitin.Tech

Lingni.Kong@Ocean University of China

Haipeng.Qu@Ocean University of China

- Beijing Chaitin Tech Co., Ltd(@ChaitinTech)

- <https://chaitin.cn/en>
- <https://realworldctf.com/>

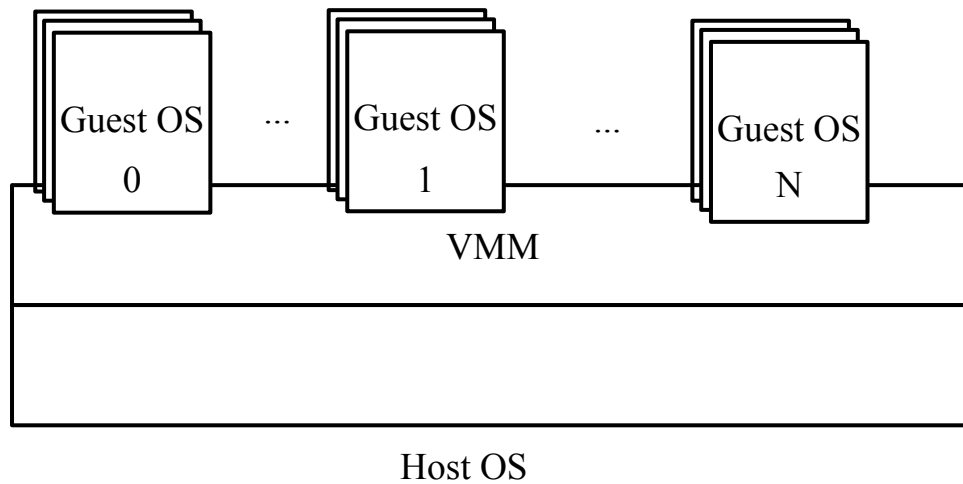


- Chaitin Security Research Lab

- Pwn2Own 2017 3rd place
- GeekPwn 2015/2016/2018/2019 awardees
 - PS4 Jailbreak, Android rooting, IoT Offensive Research, ESXi Escape
- CTF players from team b1o0p, Tea Deliverers
 - 2nd place at DEFCON 2016
 - 3rd place at DEFCON 2019
 - 1st place at HITCON 2019

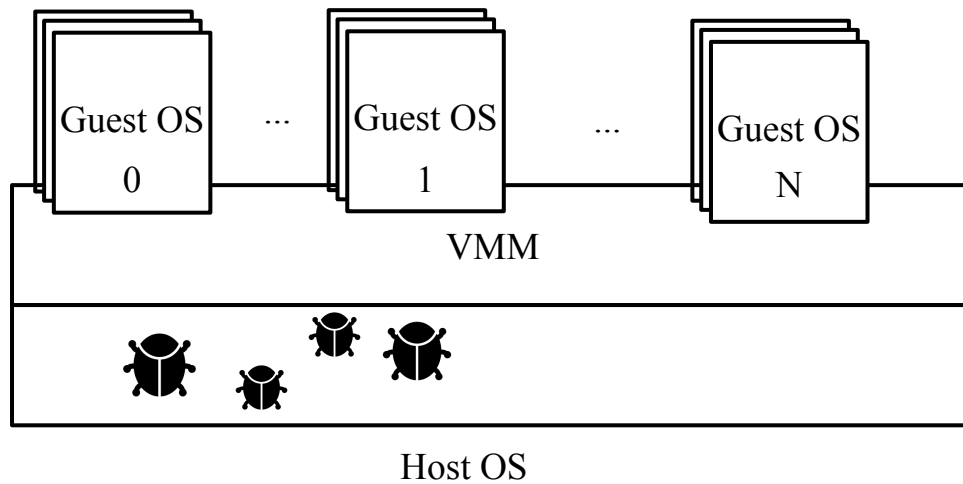
- Information Security Lab of Ocean University of China
<http://security.ouc.edu.cn/>
- OUC Security Research Lab
 - BCTF 2020 online round 1st place
 - WCTF World Hacker Masters 2019 3rd place

What is Virtual Machine Escape

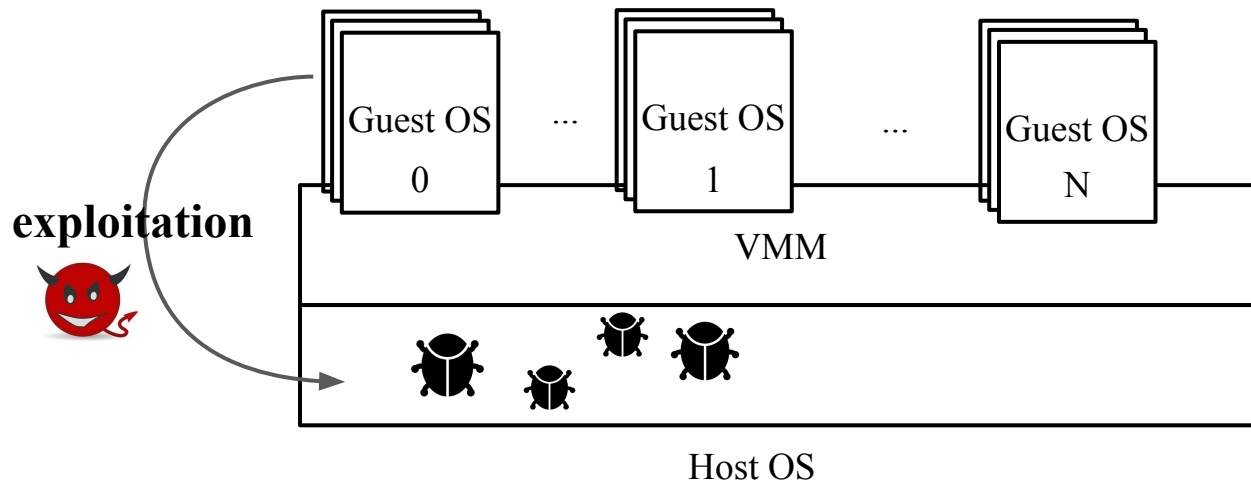


Normally, all of the sensitive behaviors of guest OS will be sanitized by the hypervisor

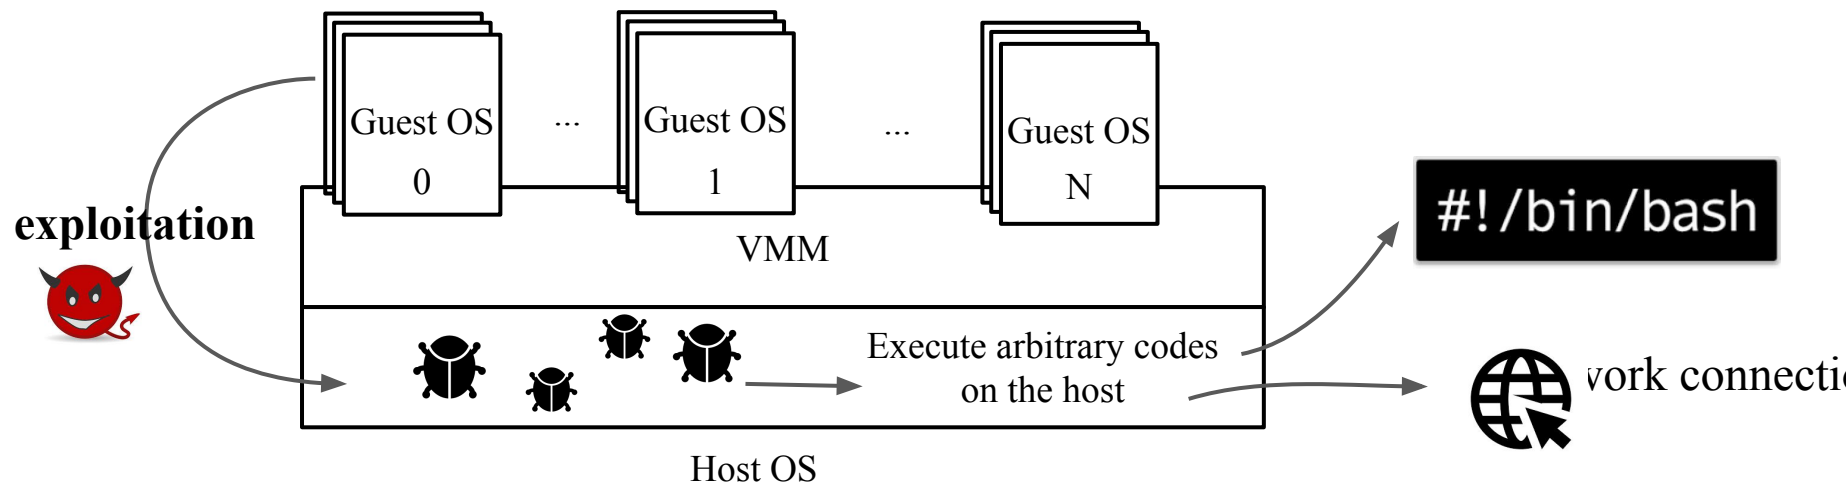
What is Virtual Machine Escape



What is Virtual Machine Escape



What is Virtual Machine Escape





Introduction of Qemu-KVM

- Open source software
- Emulator



User Space



Linux Kernel

- Kernel-based Virtual Machine
- Encapsulates VMX or SVM



User Space



Linux Kernel

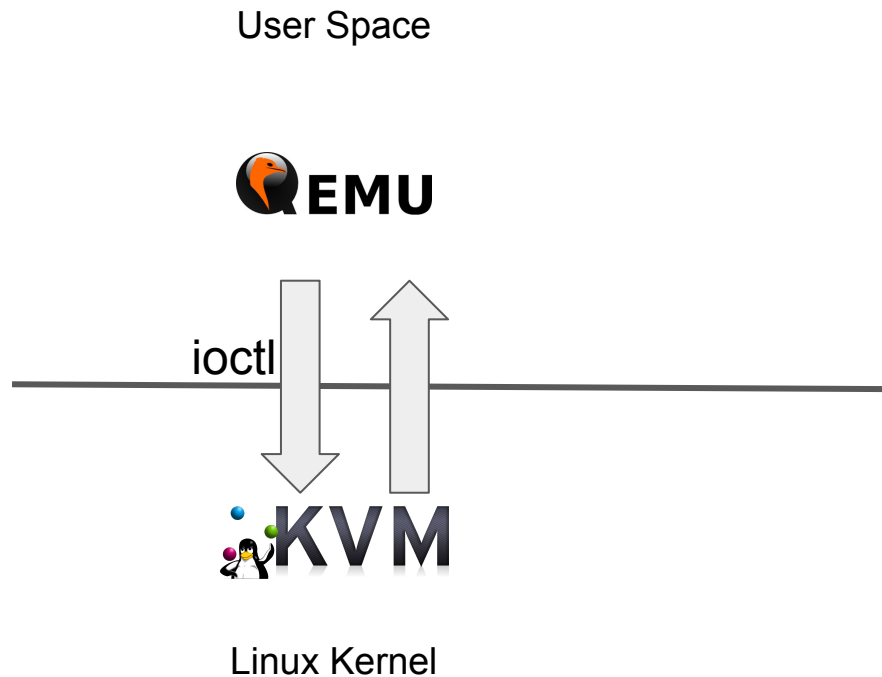
- Qemu
 - Emulates other devices
- KVM
 - Emulates CPU and memory

User Space

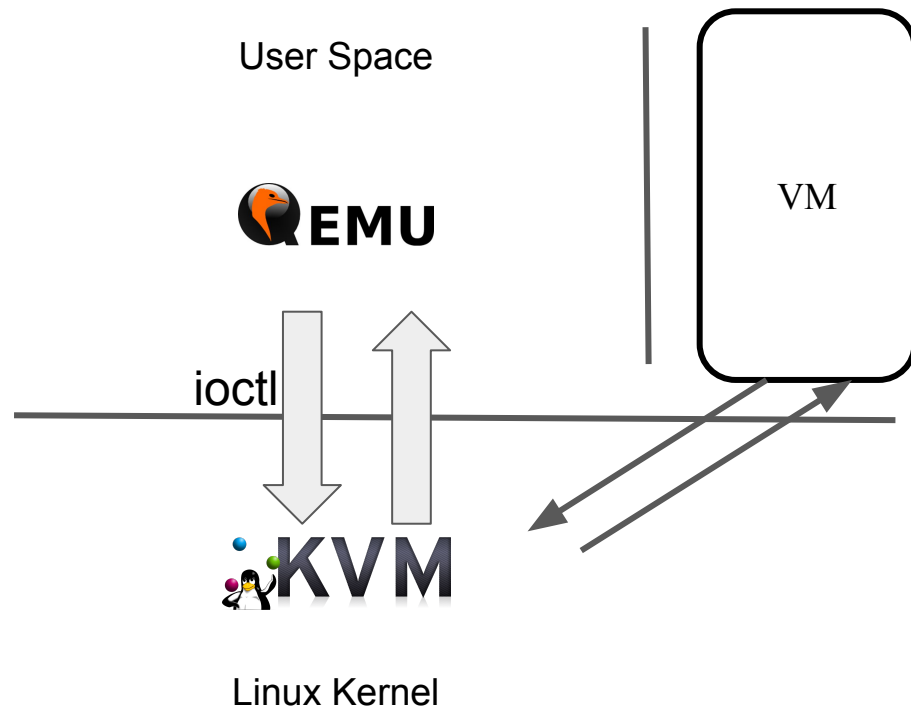


Linux Kernel

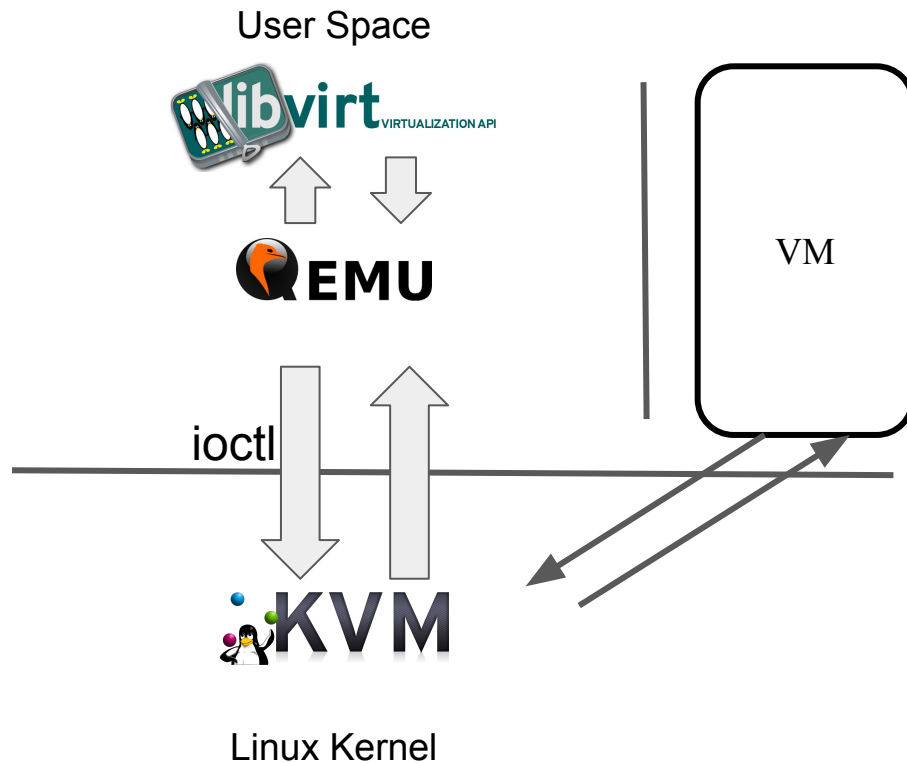
- Qemu
 - Uses ioctl and /dev/kvm
- KVM
 - Provides a series of APIs to create and run VM

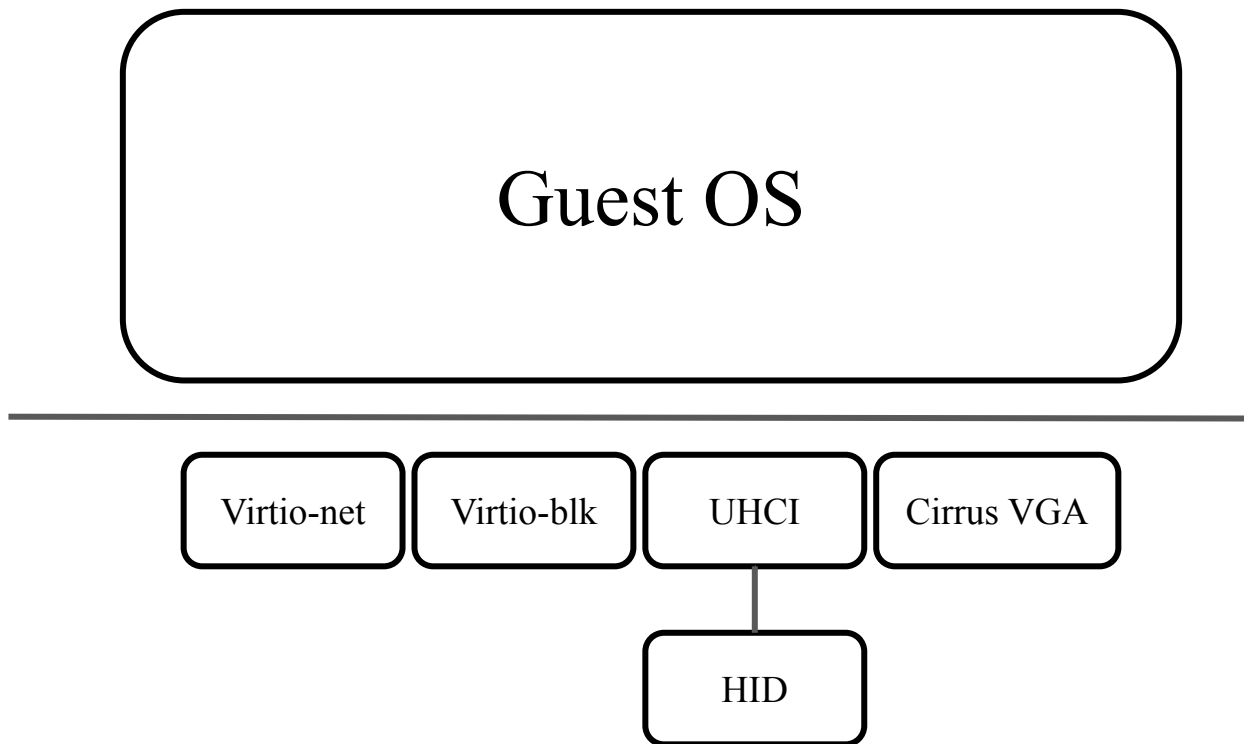


- Qemu
 - Uses ioctl and /dev/kvm
- KVM
 - Provides a series of APIs to create and run VM

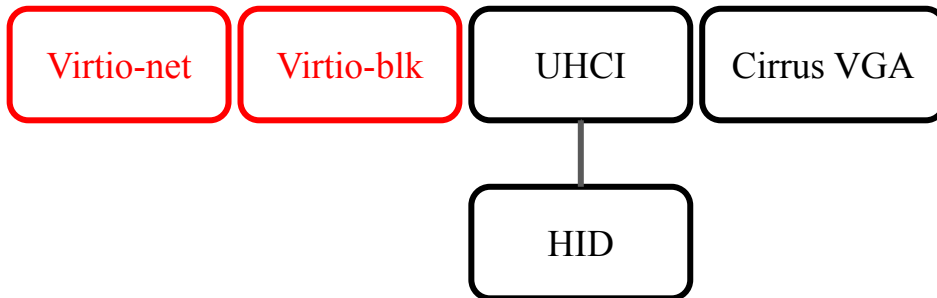


- A set of open source APIs, daemons and management tools for managing hardware virtualization
- Used by most public cloud providers.

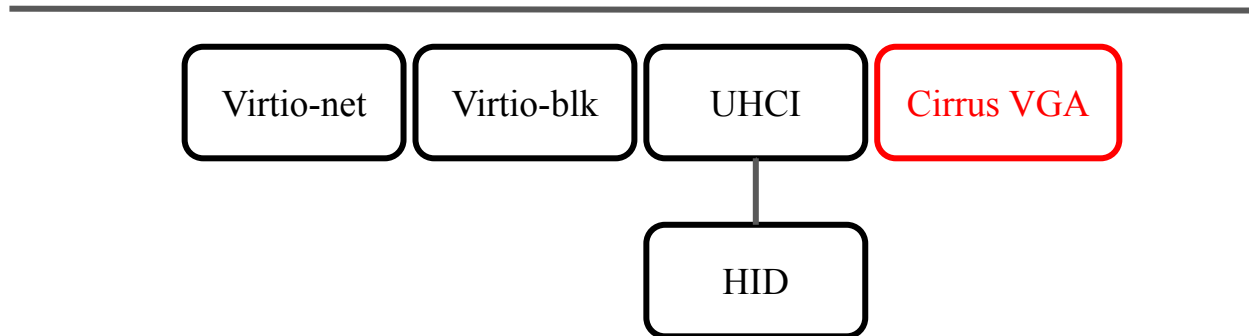




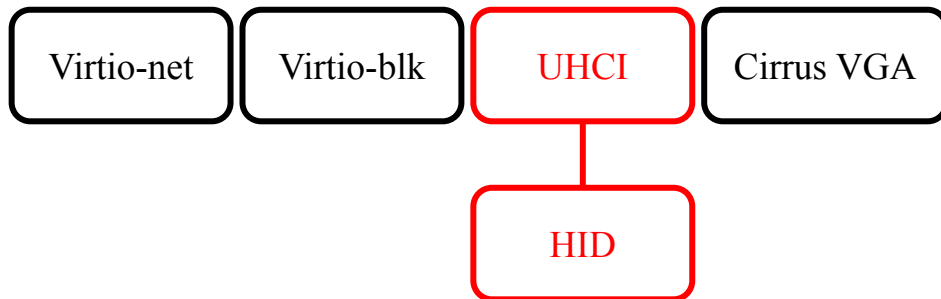
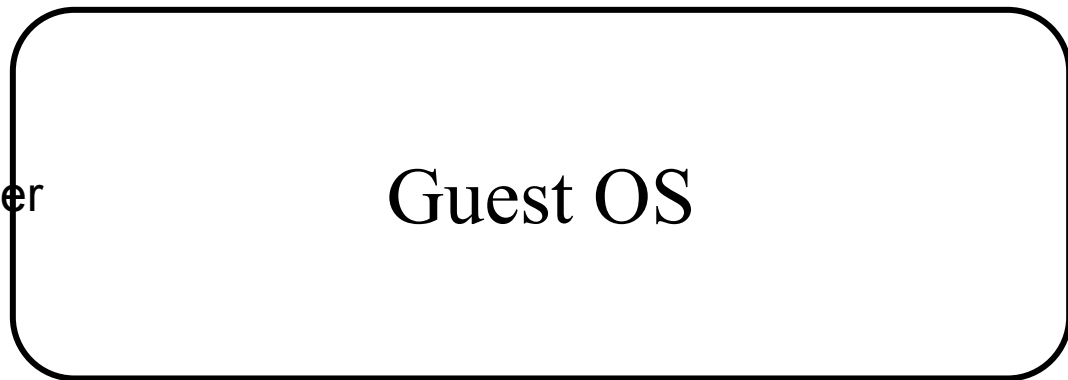
- Virtio
 - Simple
 - Few code
 - Few CVEs



- Cirrus VGA
 - Many CVEs
 - Hard to exploit

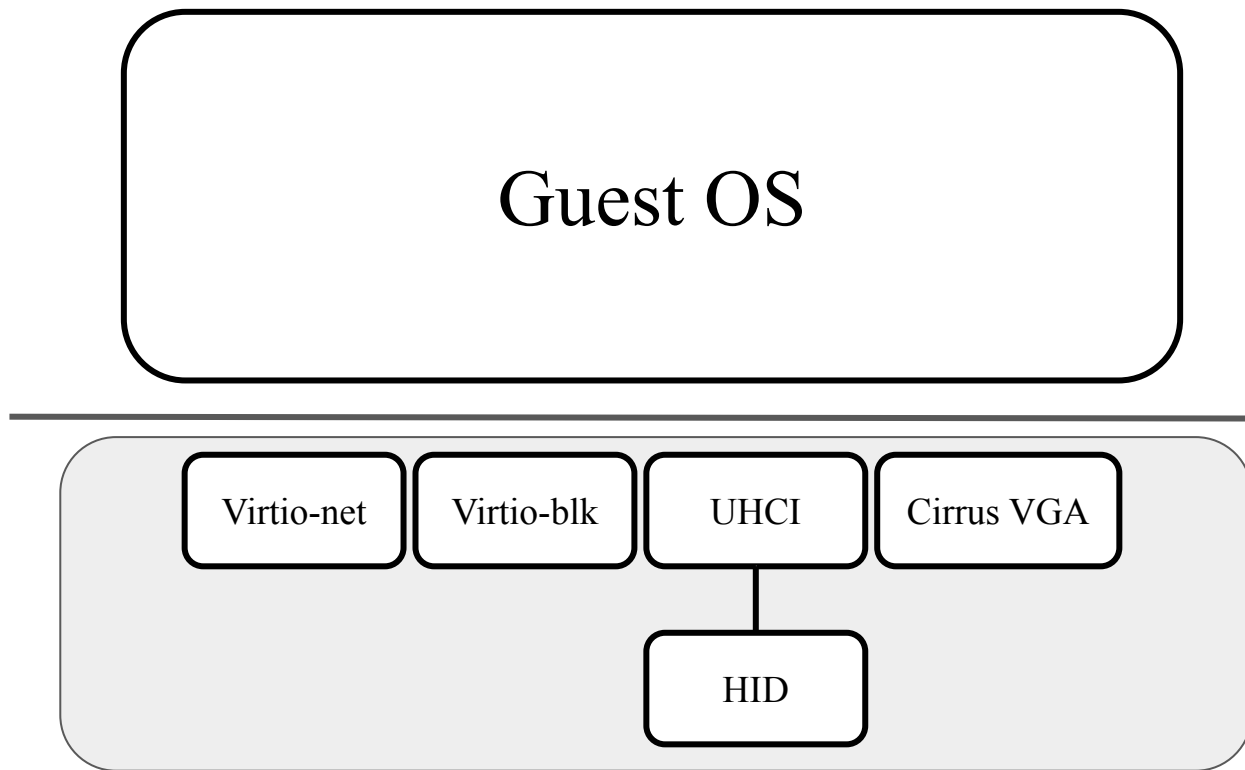


- UHCI
 - Universal Host Controller Interface
 - USB 1.0
- HID
 - Human Interface Device
 - mouse/keyboard



Why hard to escape from the public cloud?

- Lack of good vulnerabilities
- Lack of further information





CVE-2020-14364

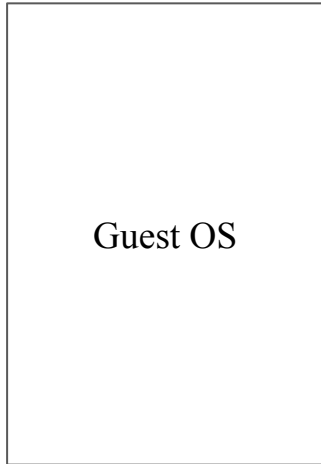
- Reported at 2020.8.13
- Redhat fixed it and disclosed it at 2020.8.24

How Does Guest OS send usb packets ?



Universal Host Controller Interface

registers



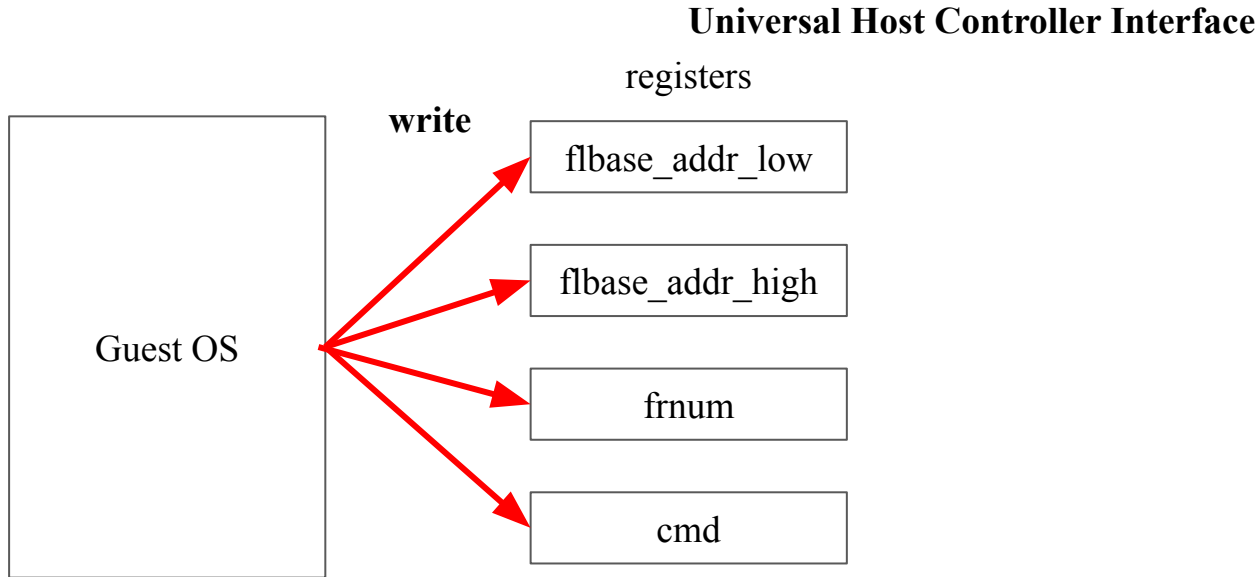
flbase_addr_low

flbase_addr_high

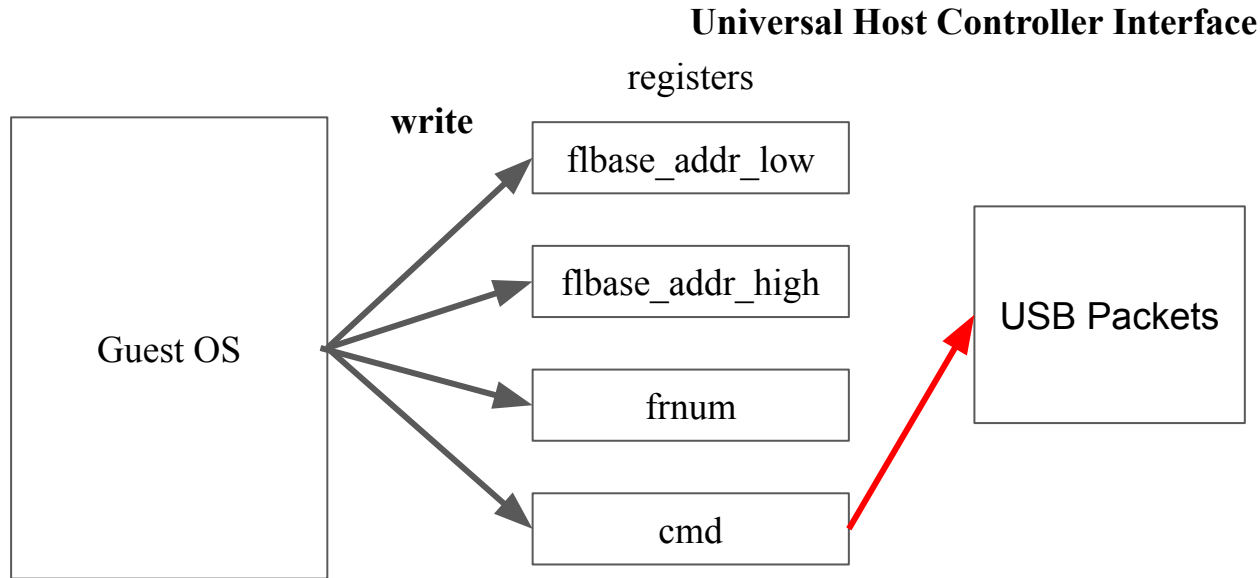
frnum

cmd

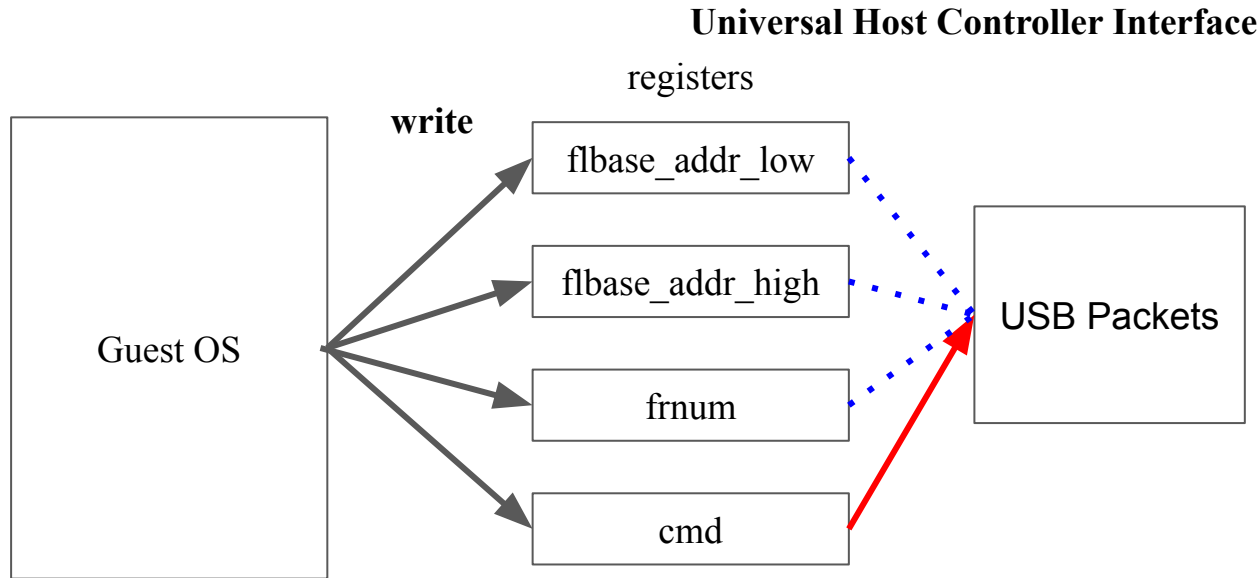
How Does Guest OS send usb packets ?



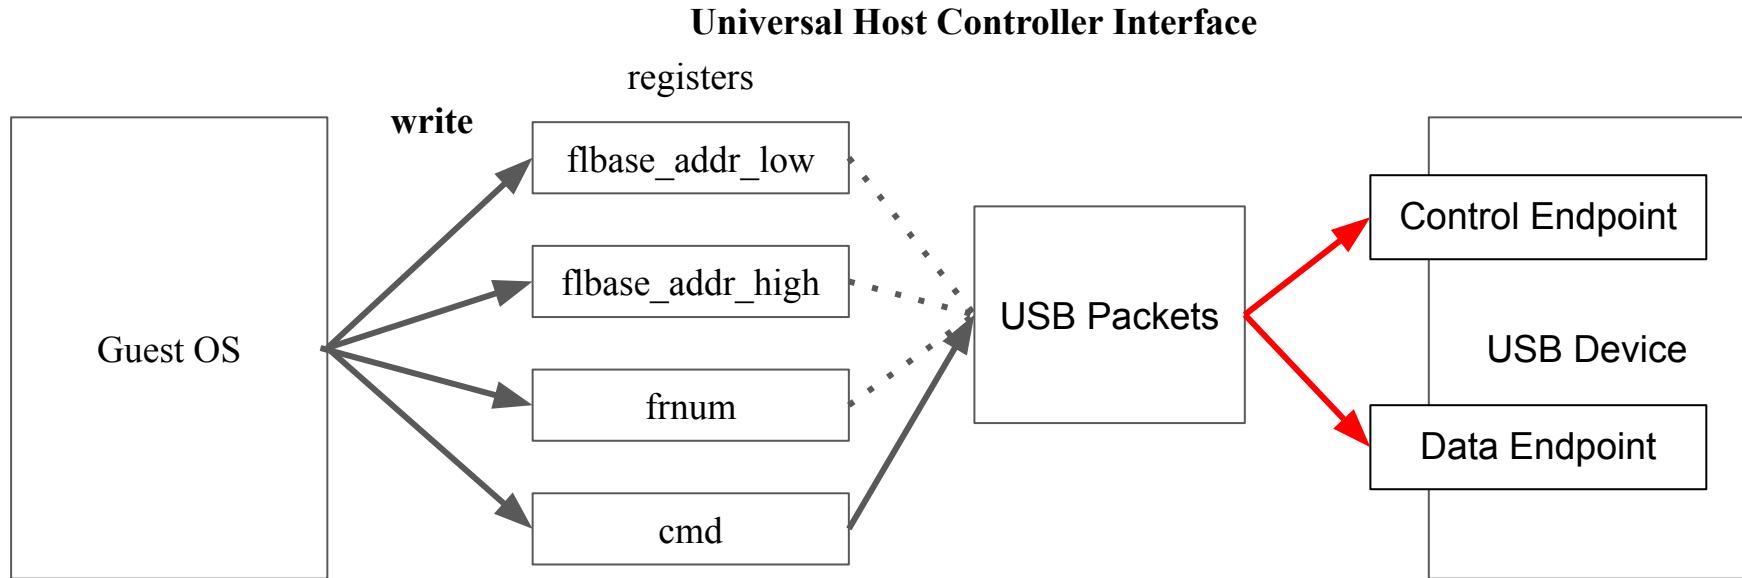
How Does Guest OS send usb packets ?



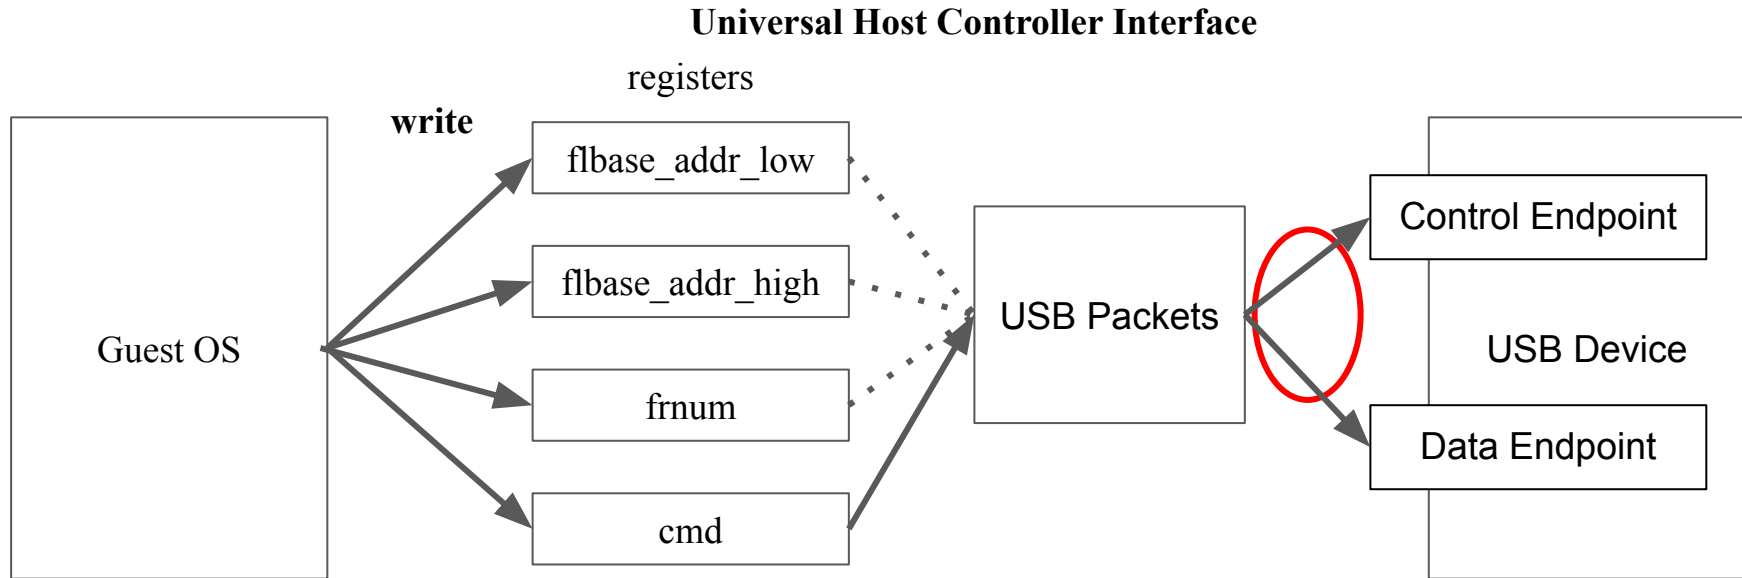
How Does Guest OS send usb packets ?



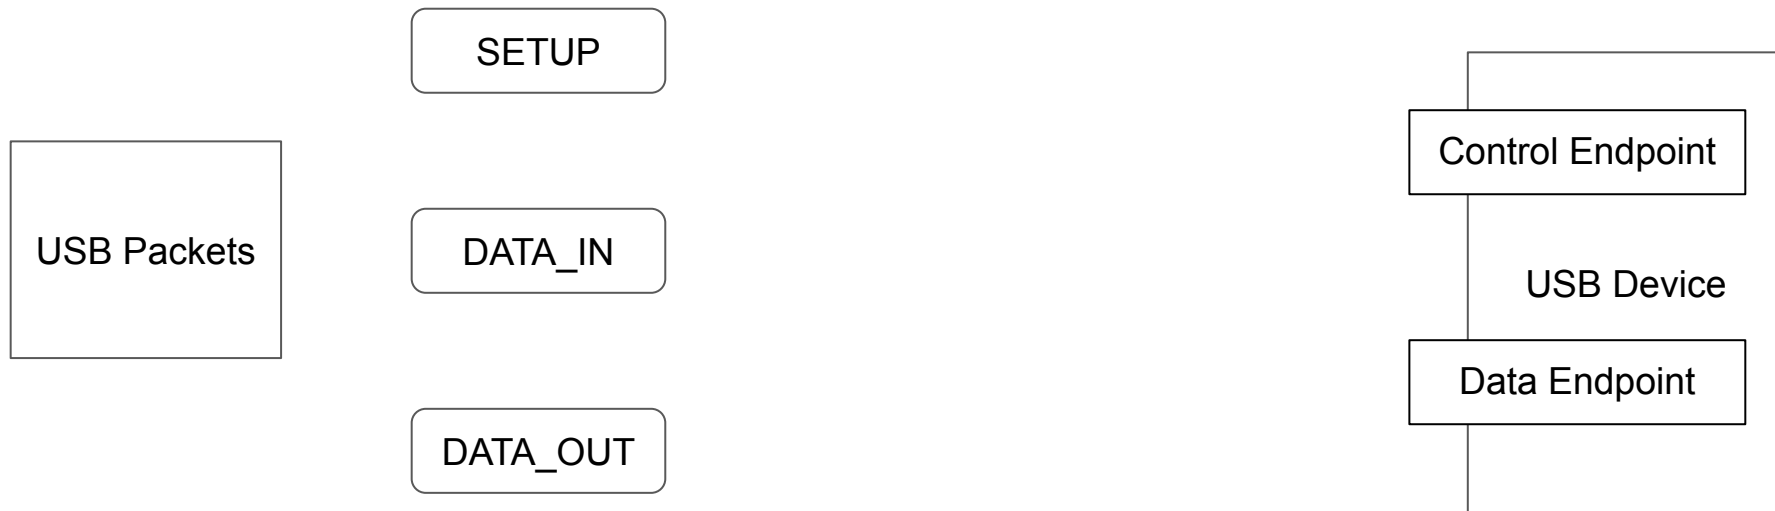
How Does Guest OS send usb packets ?



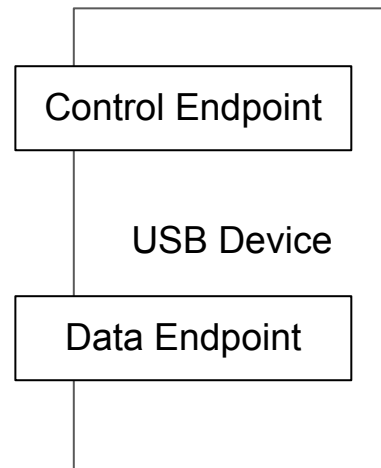
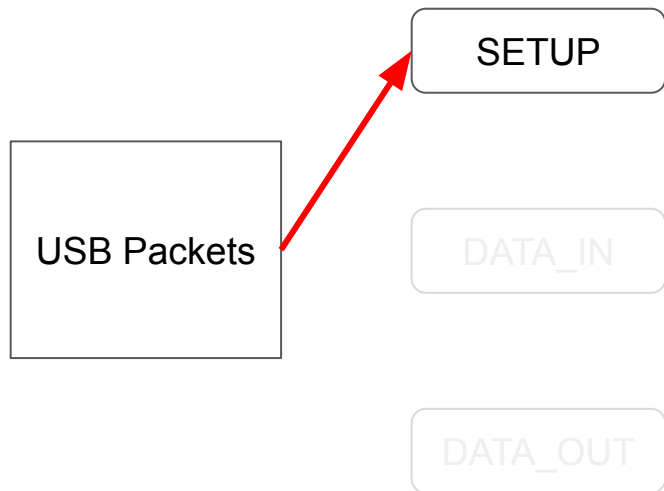
How Does Guest OS send usb packets ?



How Do Qemu transfer USB packets?



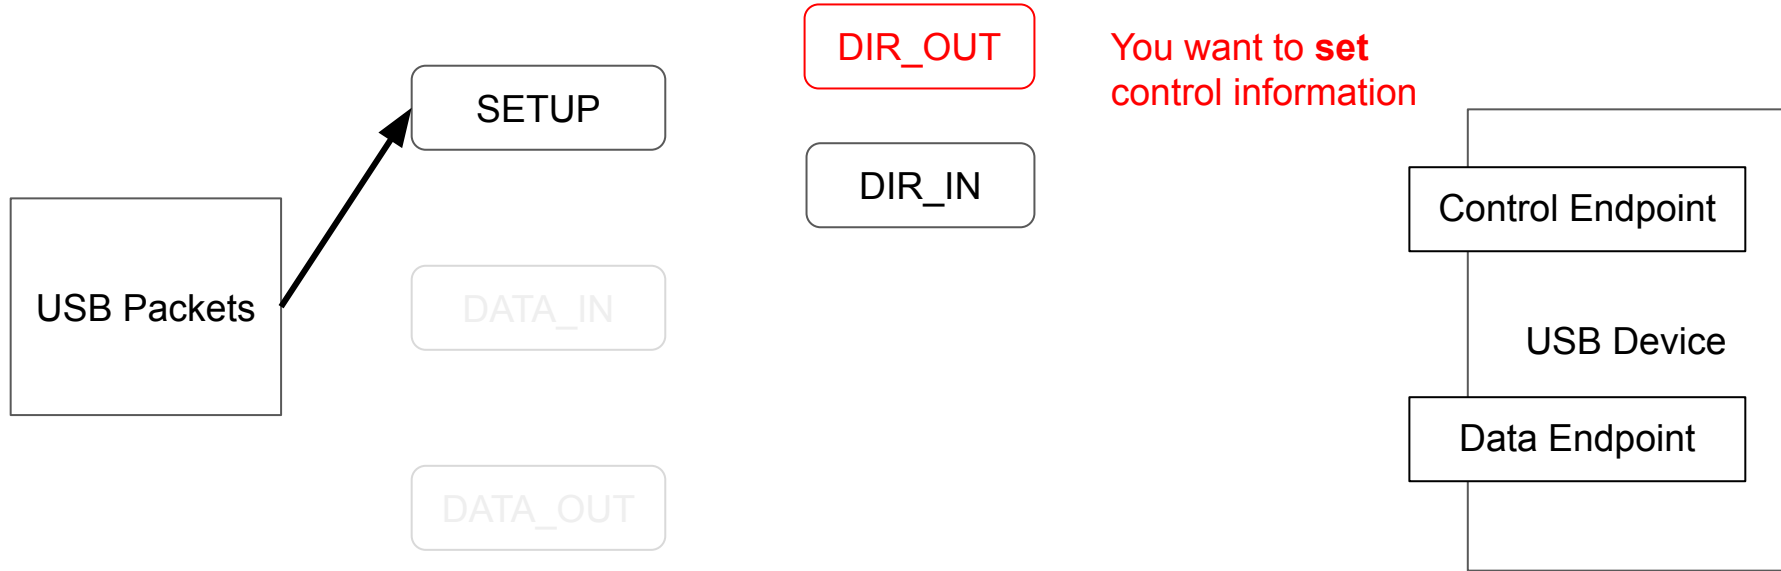
How Do Qemu transfer USB packets?



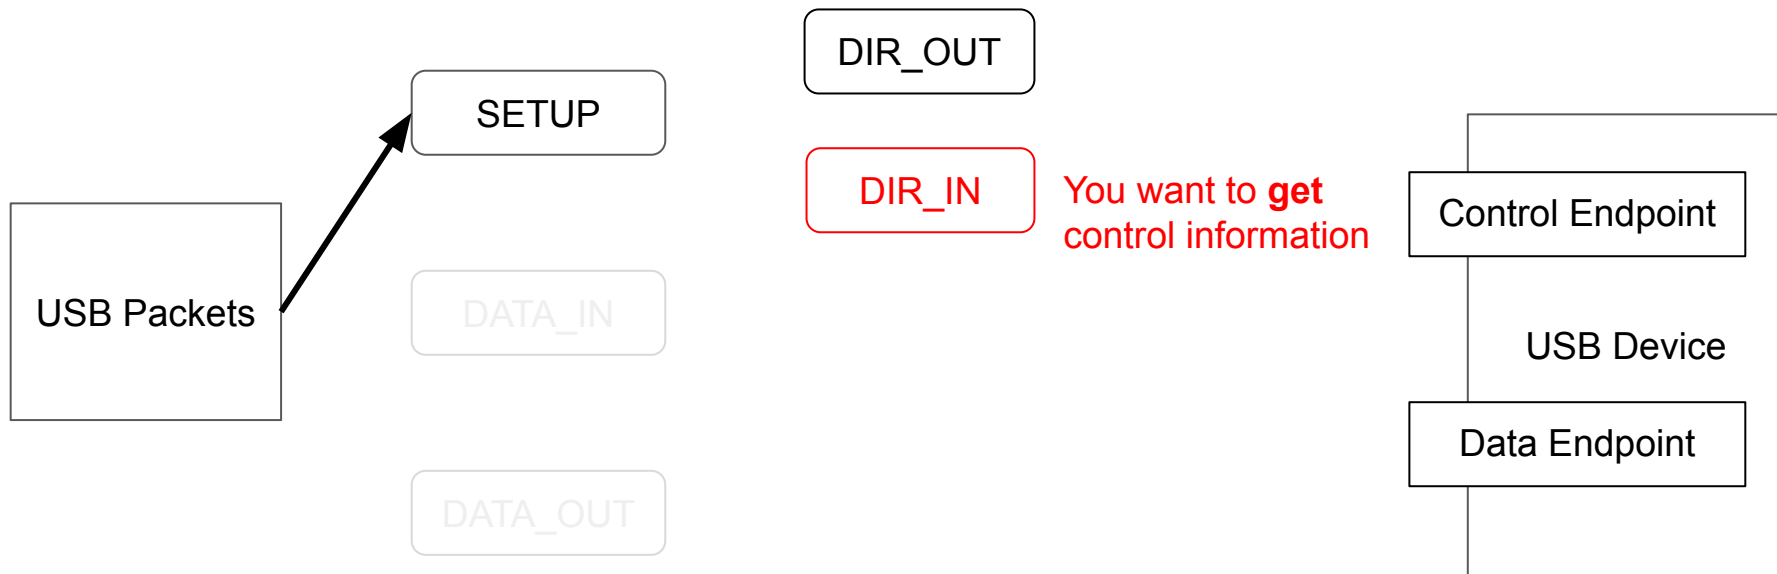
How Do Qemu transfer USB packets?



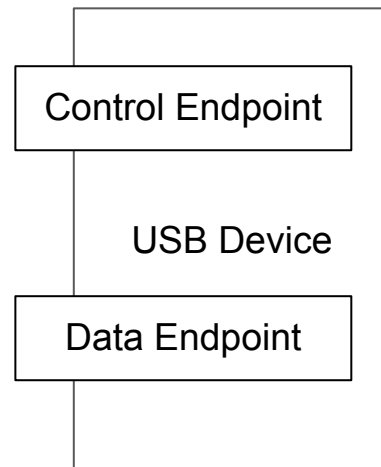
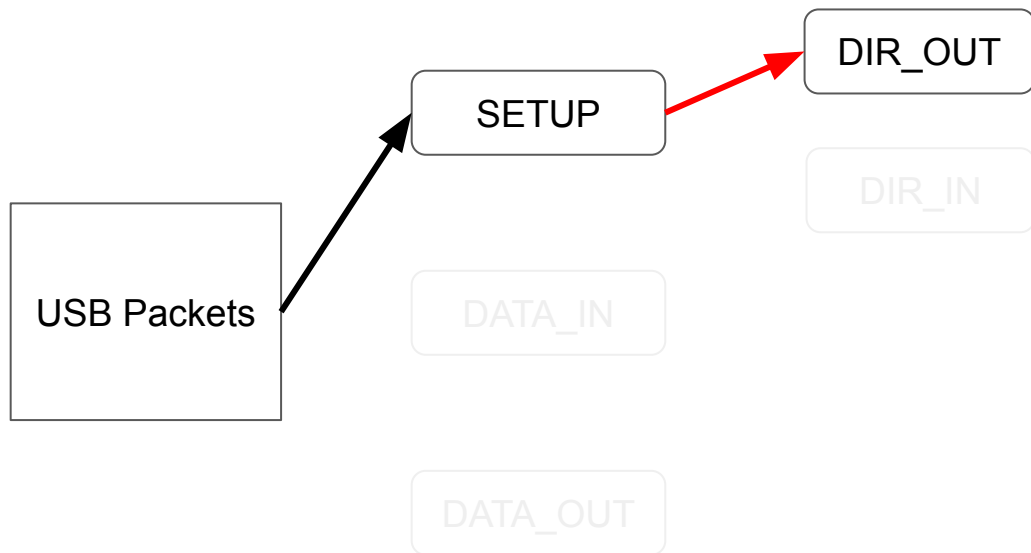
How Do Qemu transfer USB packets?



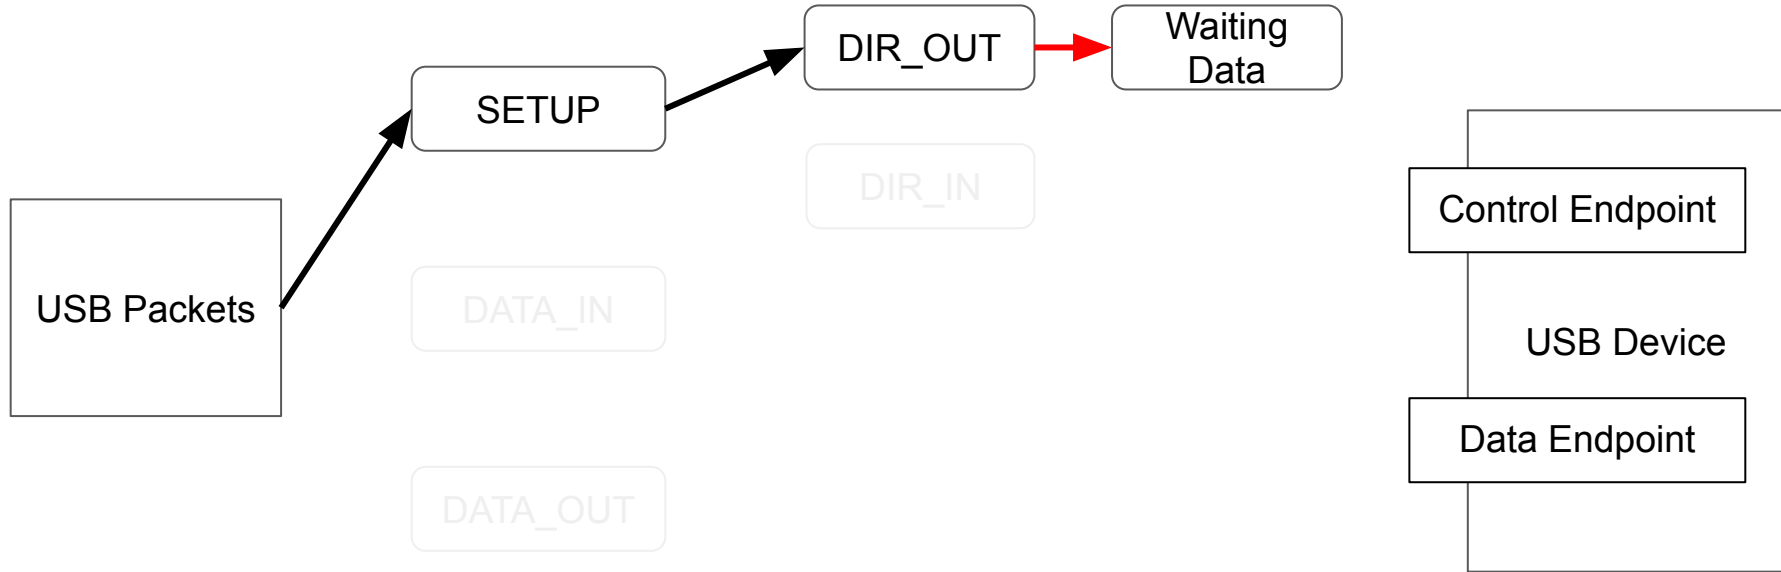
How Do Qemu transfer USB packets?



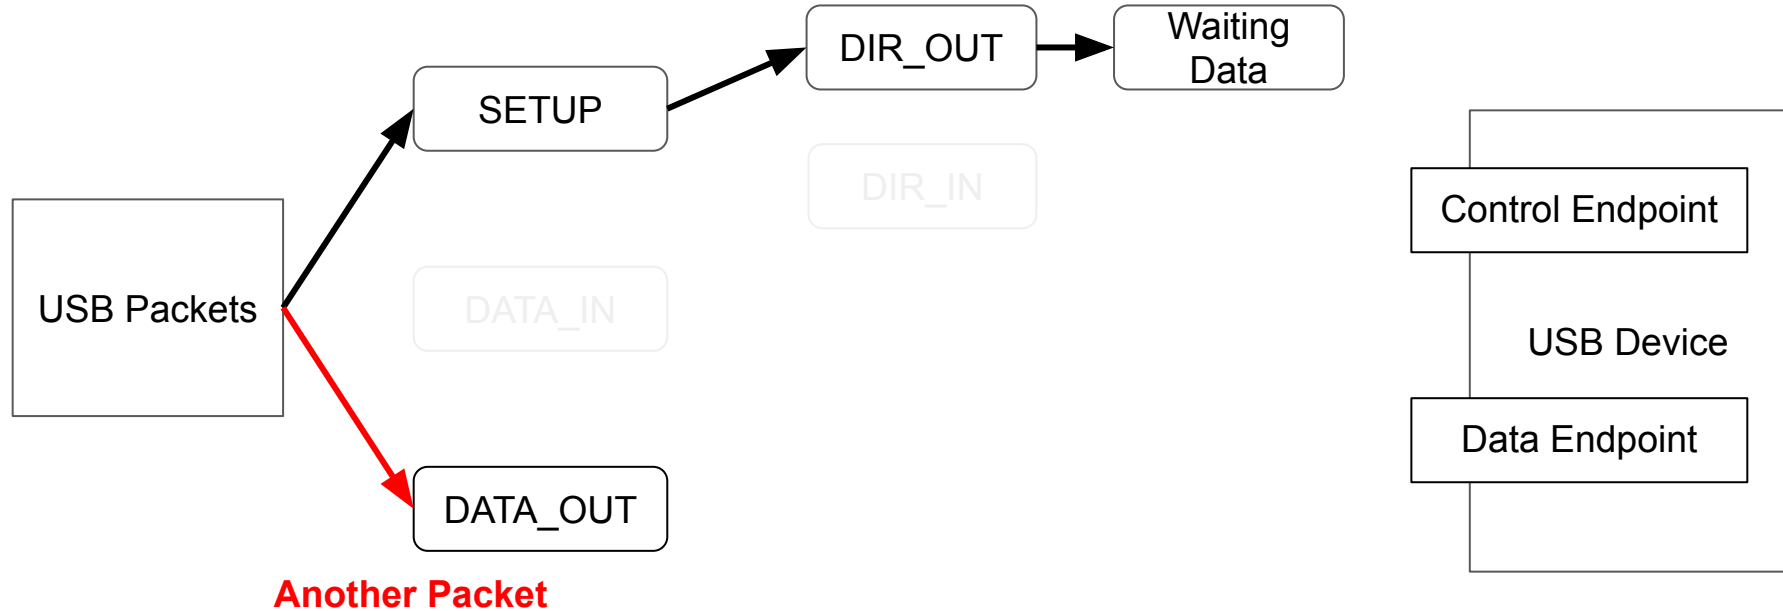
How Do Qemu transfer USB packets?



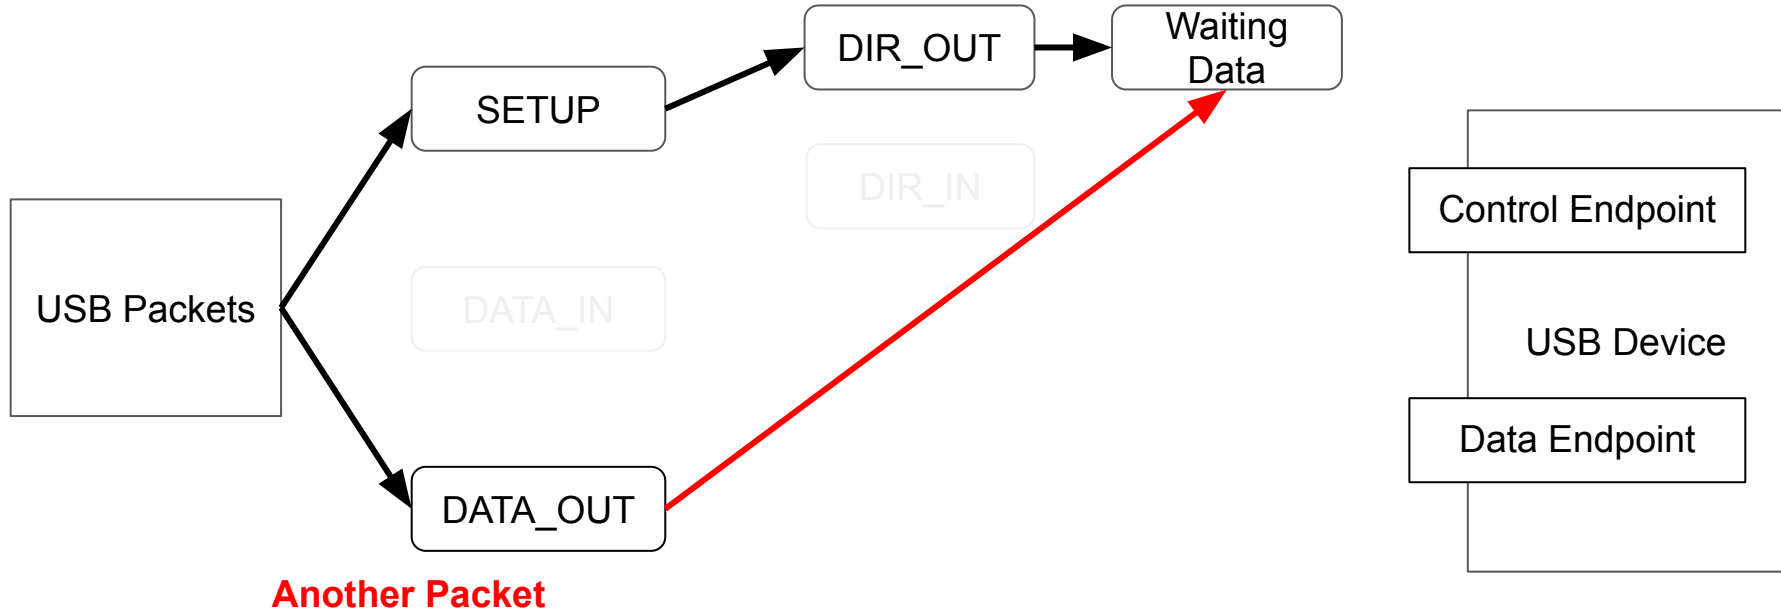
How Do Qemu transfer USB packets?



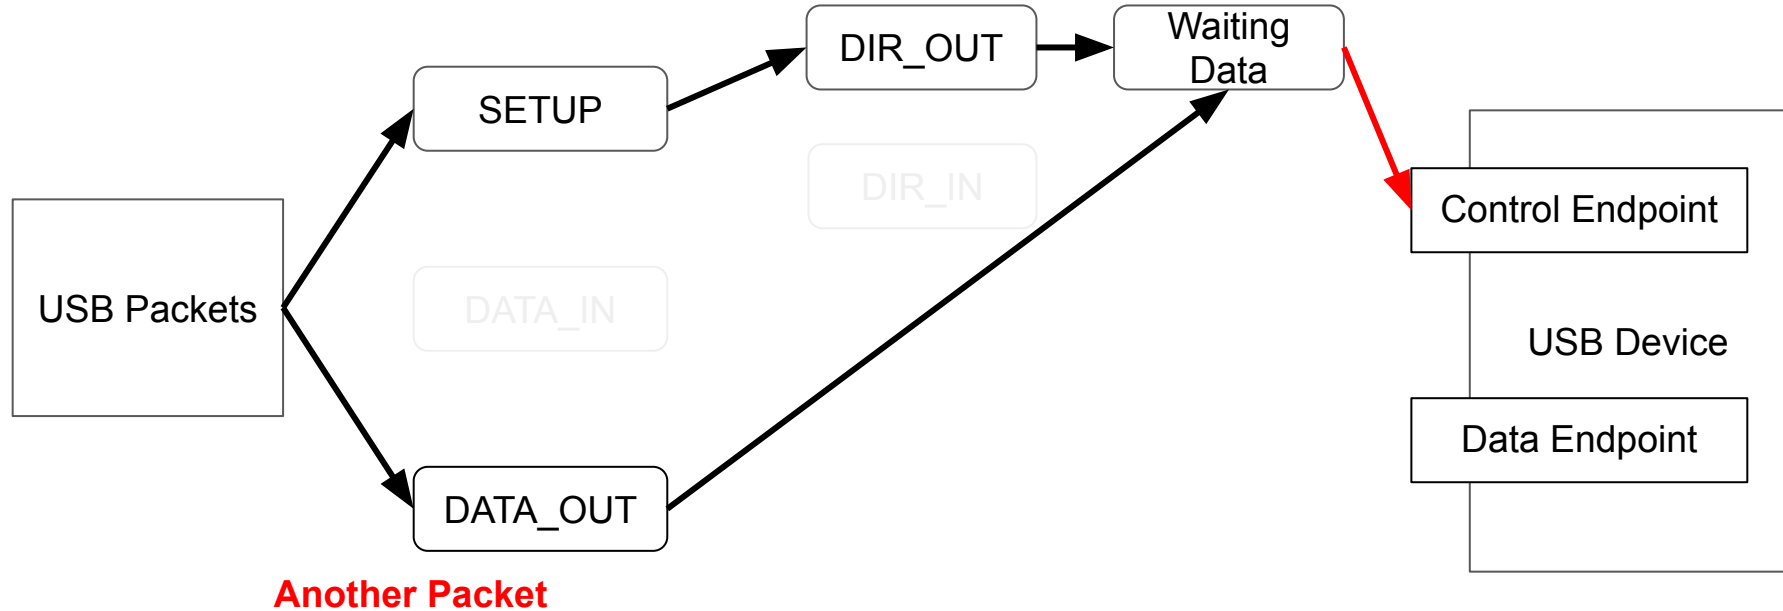
How Do Qemu transfer USB packets?



How Do Qemu transfer USB packets?



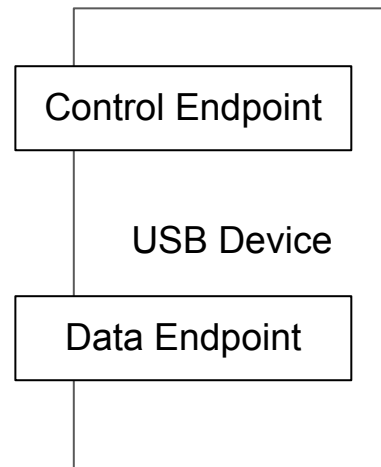
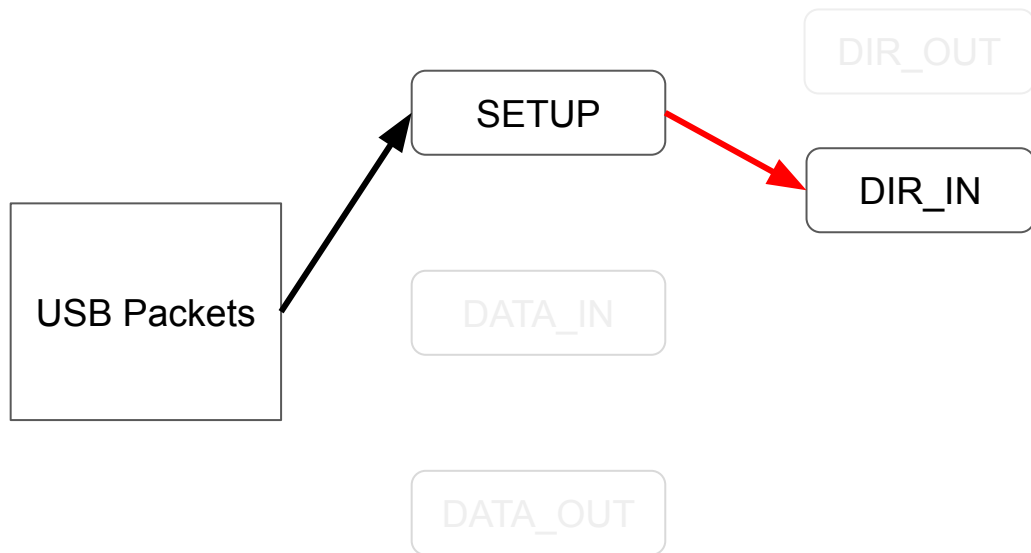
How Do Qemu transfer USB packets?



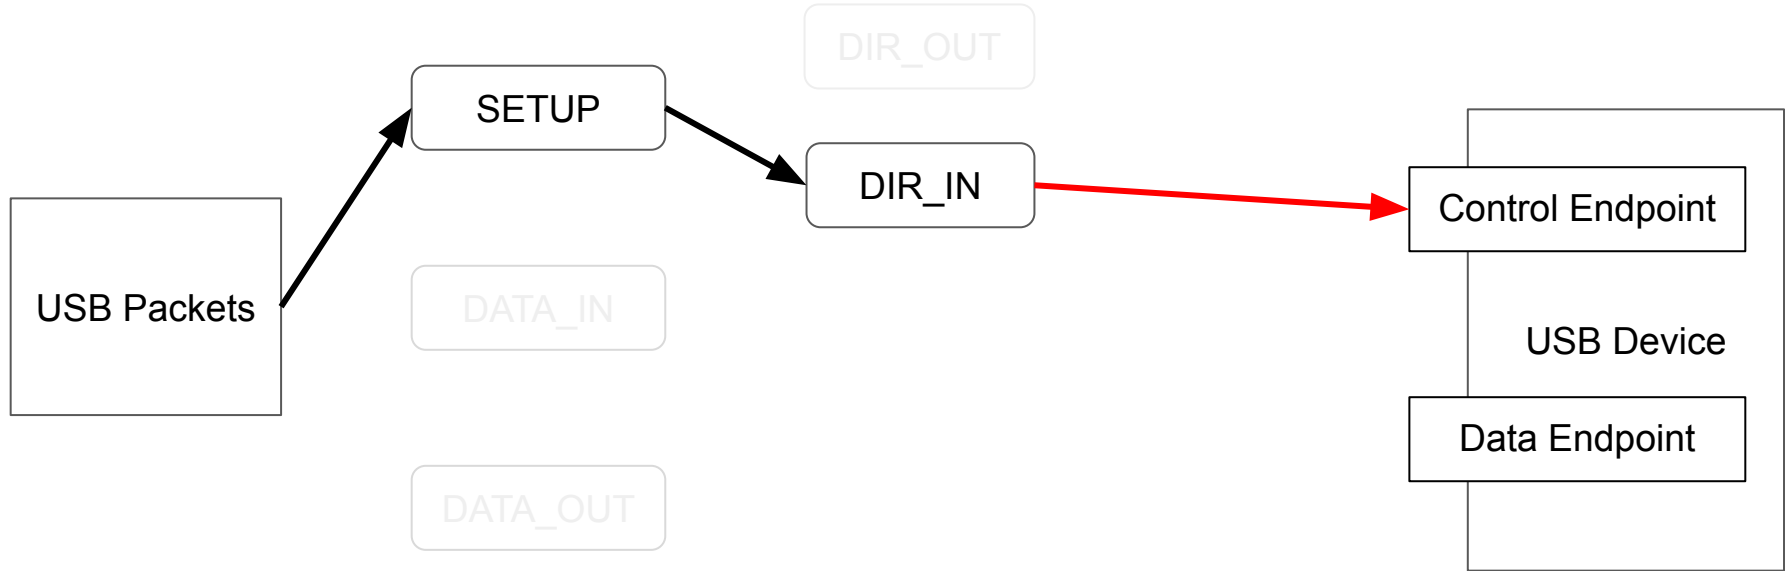
How Do Qemu transfer USB packets?



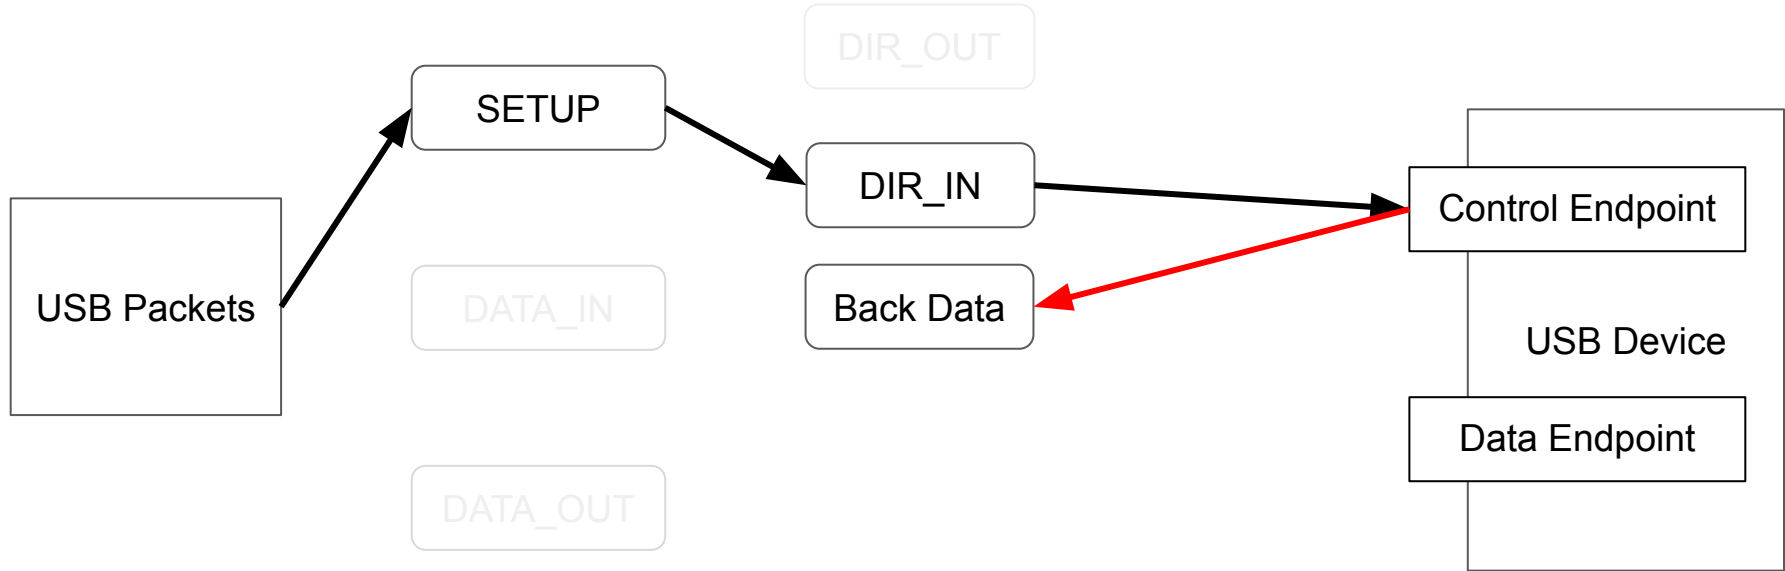
How Do Qemu transfer USB packets?



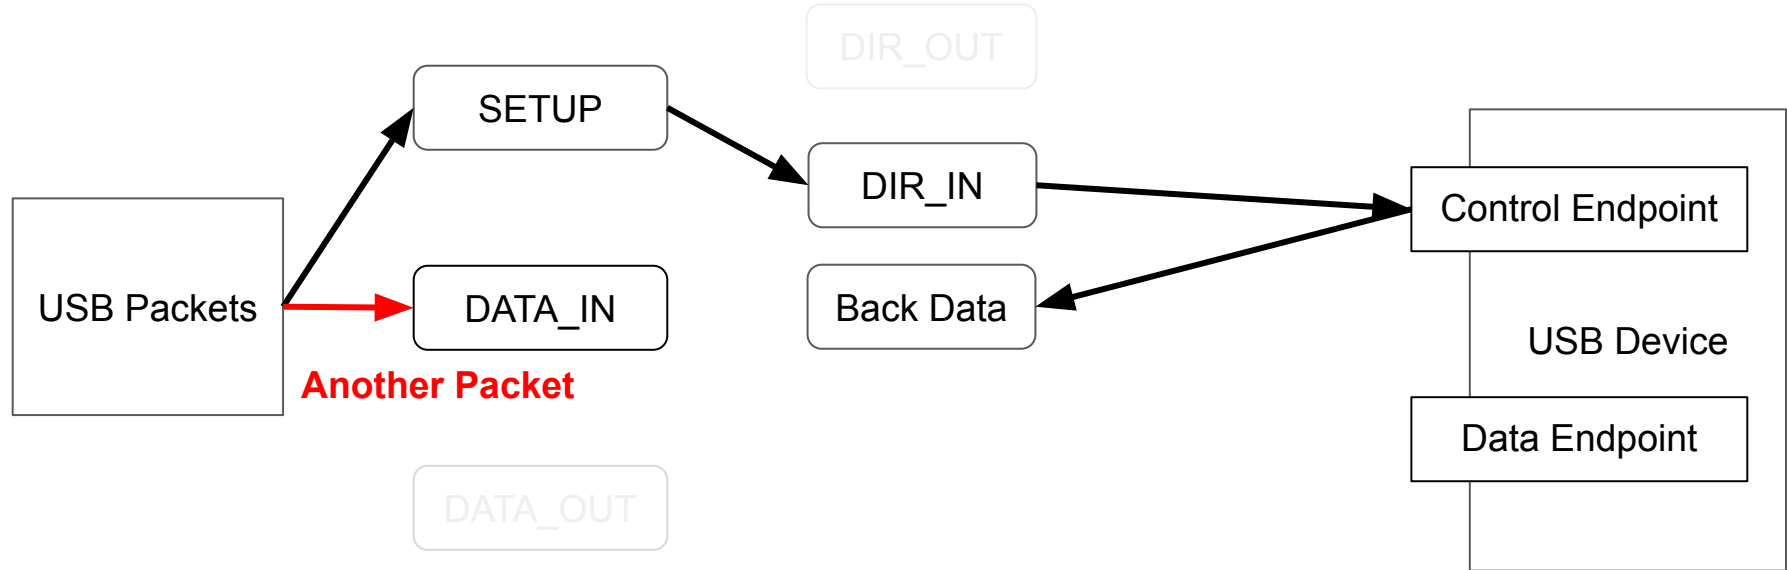
How Do Qemu transfer USB packets?



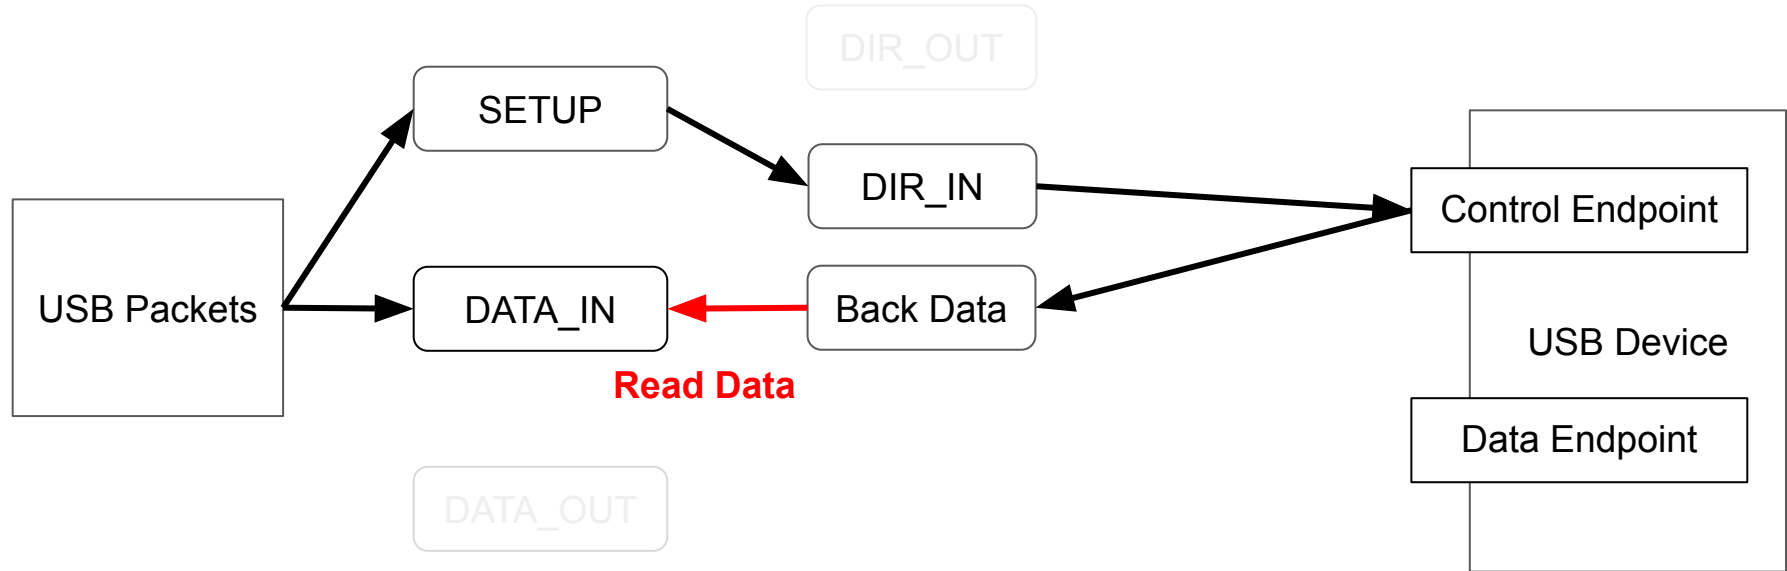
How Do Qemu transfer USB packets?



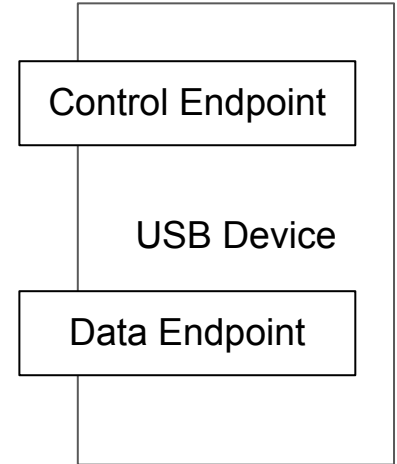
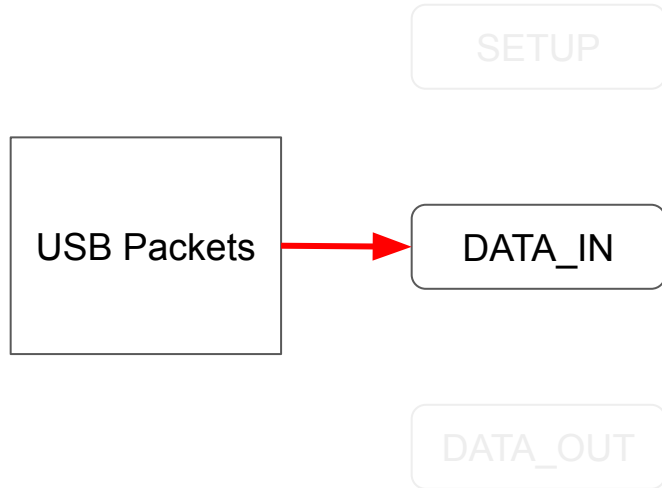
How Do Qemu transfer USB packets?



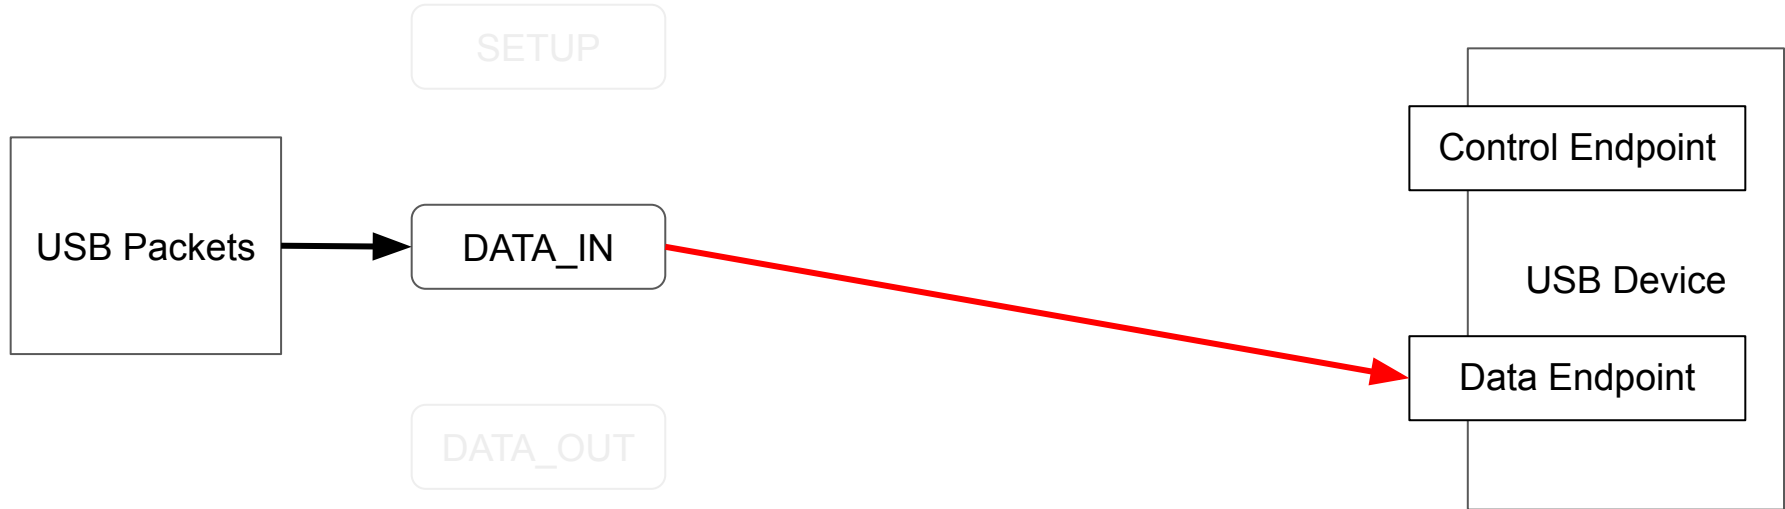
How Qemu transfer USB packets?



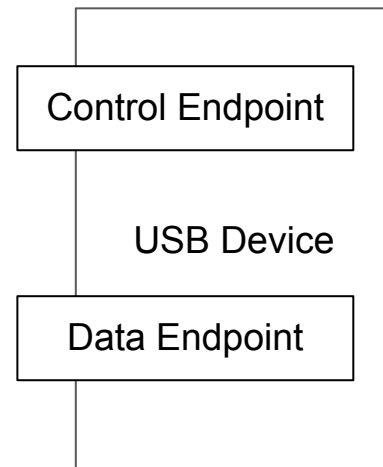
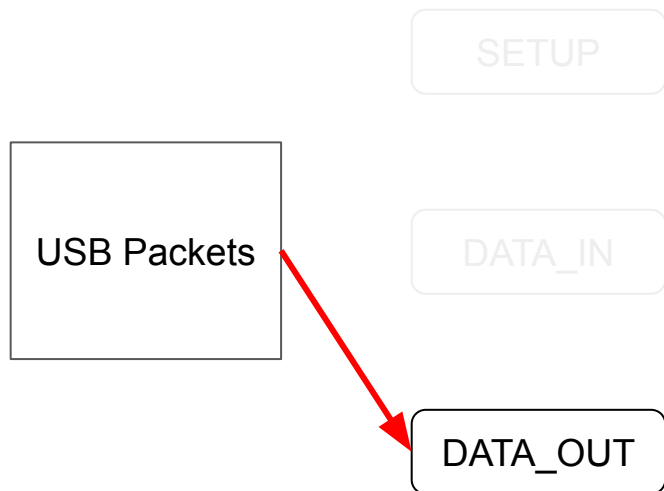
How Do Qemu transfer USB packets?



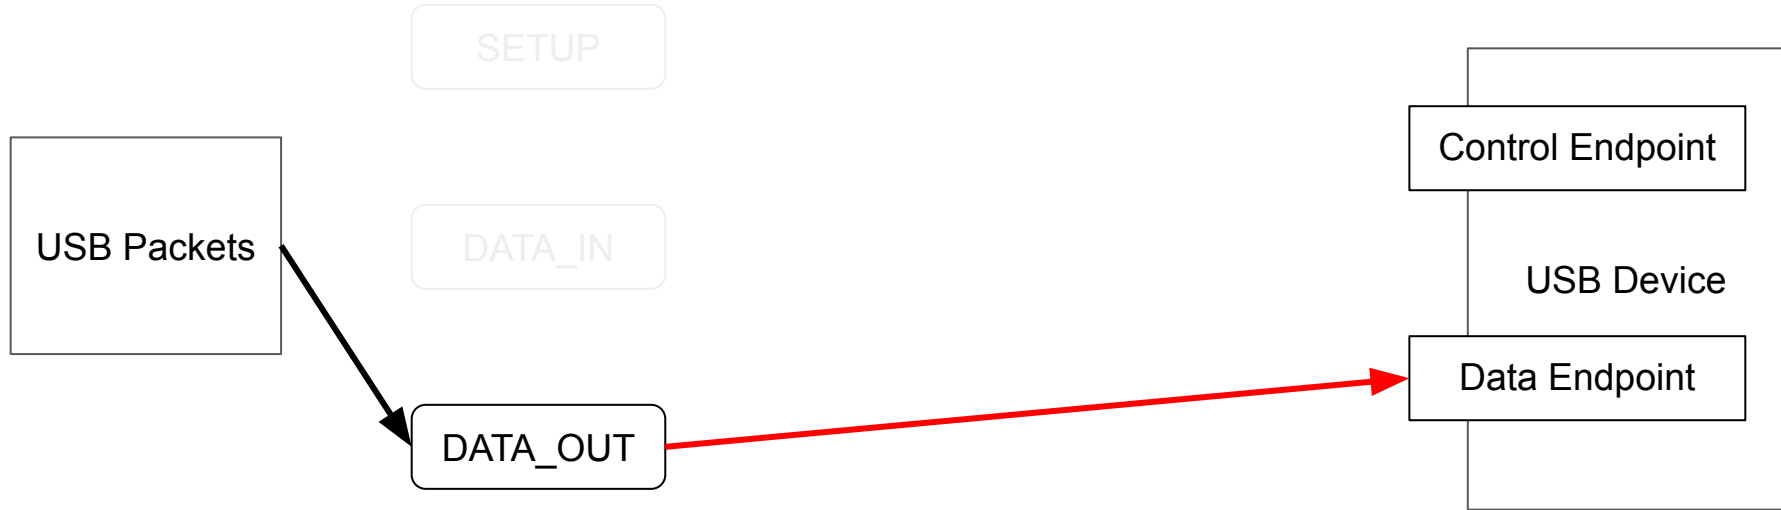
How Do Qemu transfer USB packets?



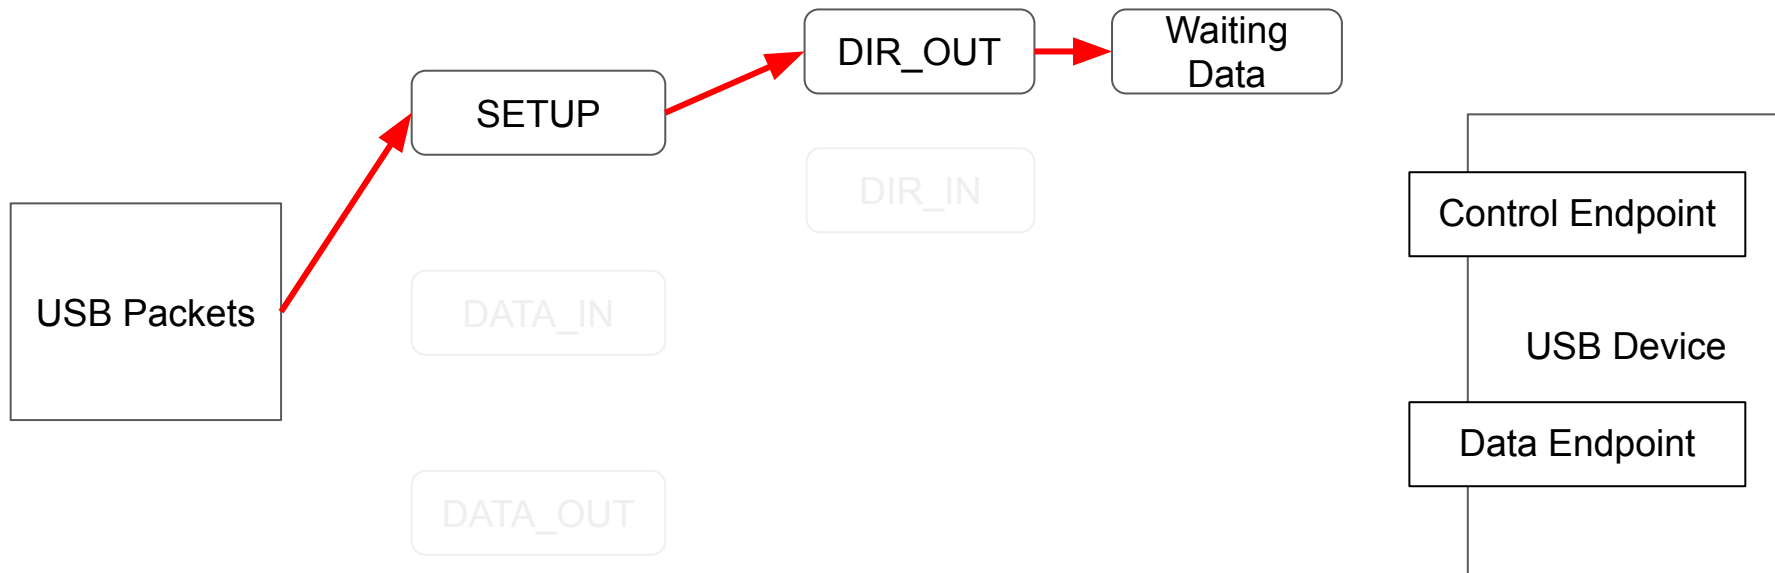
How Do Qemu transfer USB packets?

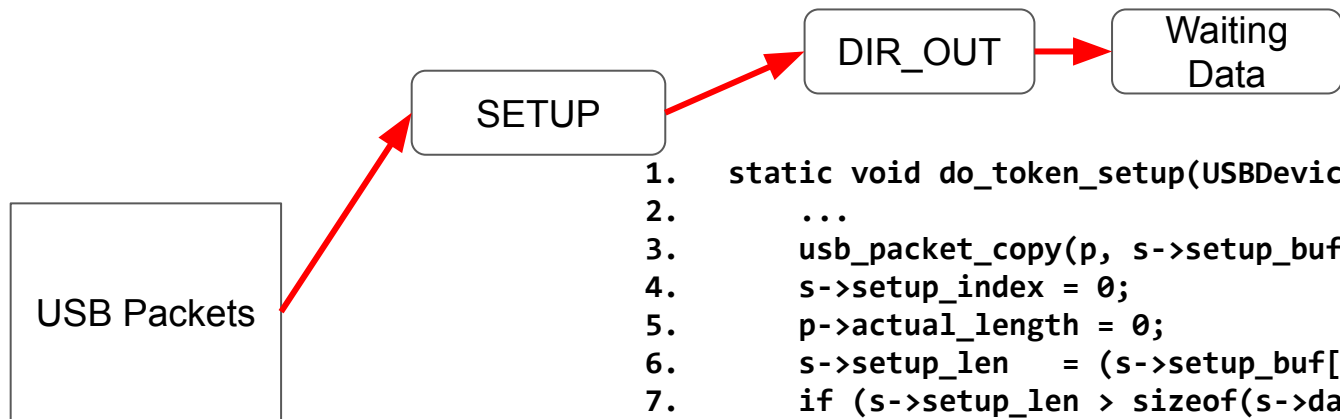


How Do Qemu transfer USB packets?

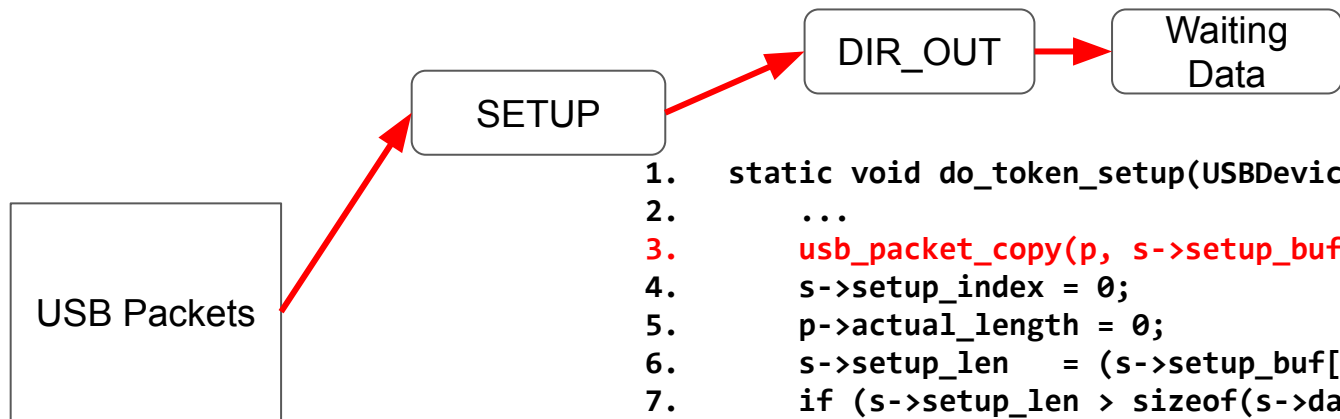


CVE-2020-14364

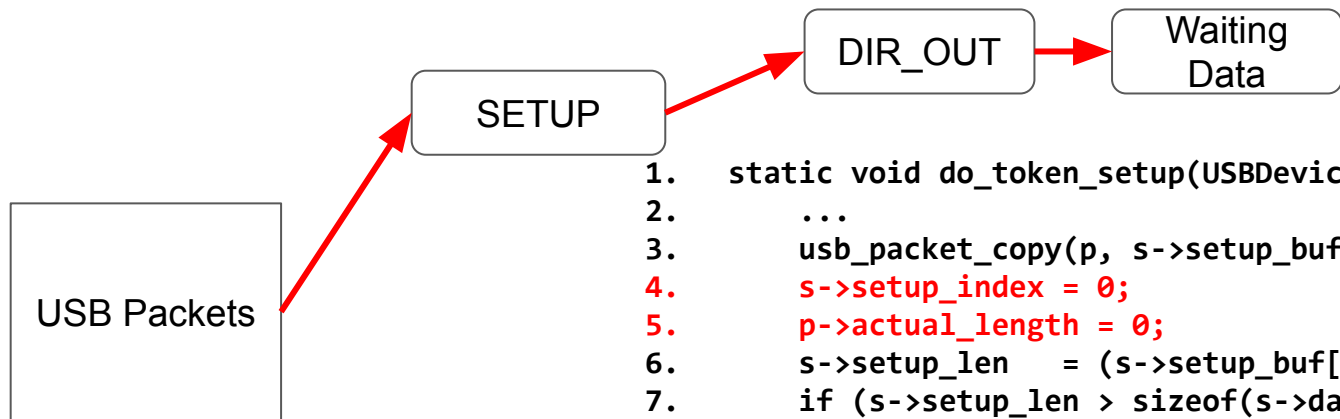




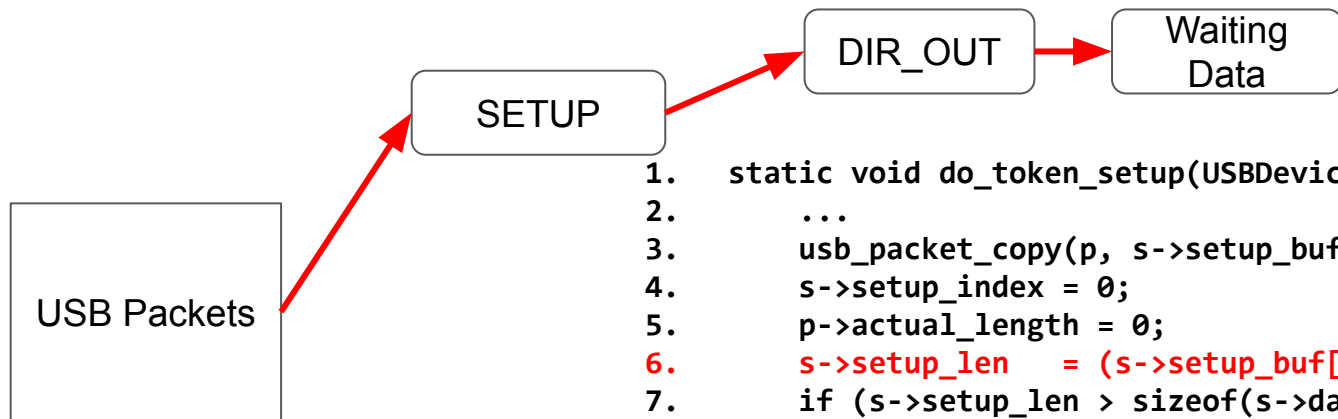
```
1. static void do_token_setup(USBDevice *s, USBPacket *p){
2.     ...
3.     usb_packet_copy(p, s->setup_buf, p->iov.size);
4.     s->setup_index = 0;
5.     p->actual_length = 0;
6.     s->setup_len = (s->setup_buf[7] << 8) | s->setup_buf[6];
7.     if (s->setup_len > sizeof(s->data_buf)) {
8.         fprintf(stderr,
9.             "usb_generic_handle_packet: ctrl buffer too
small (%d > %zu)\n",
10.            s->setup_len, sizeof(s->data_buf));
11.         p->status = USB_RET_STALL;
12.         return;
13.     }
14.     if (s->setup_buf[0] & USB_DIR_IN) {
15.         ...
16.     }else{
17.         s->setup_state = SETUP_STATE_DATA;
```



```
1. static void do_token_setup(USBDevice *s, USBPacket *p){
2.     ...
3.     usb_packet_copy(p, s->setup_buf, p->iov.size);
4.     s->setup_index = 0;
5.     p->actual_length = 0;
6.     s->setup_len = (s->setup_buf[7] << 8) | s->setup_buf[6];
7.     if (s->setup_len > sizeof(s->data_buf)) {
8.         fprintf(stderr,
9.             "usb_generic_handle_packet: ctrl buffer too
small (%d > %zu)\n",
10.            s->setup_len, sizeof(s->data_buf));
11.         p->status = USB_RET_STALL;
12.         return;
13.     }
14.     if (s->setup_buf[0] & USB_DIR_IN) {
15.         ...
16.     }else{
17.         s->setup_state = SETUP_STATE_DATA;
```



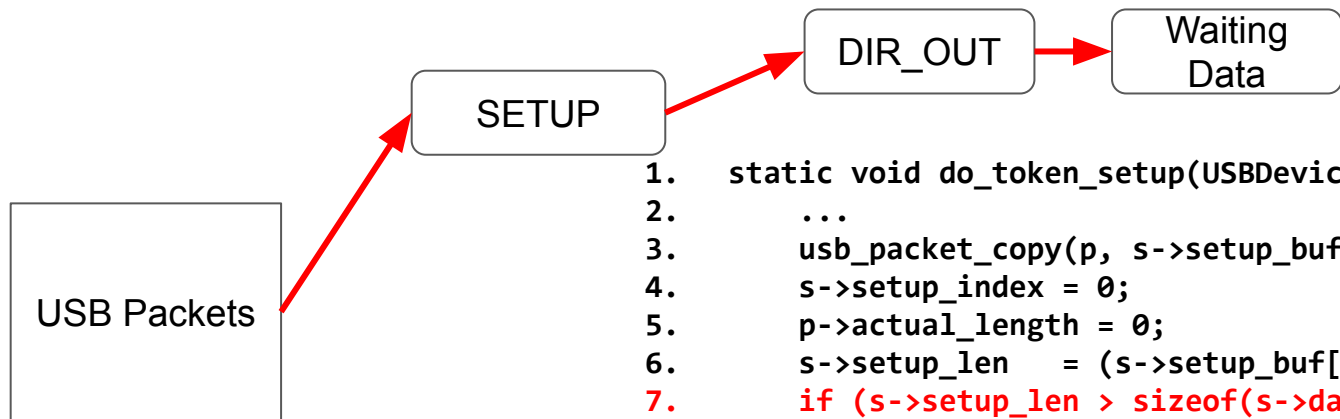
```
1. static void do_token_setup(USBDevice *s, USBPacket *p){
2.     ...
3.     usb_packet_copy(p, s->setup_buf, p->iov.size);
4.     s->setup_index = 0;
5.     p->actual_length = 0;
6.     s->setup_len = (s->setup_buf[7] << 8) | s->setup_buf[6];
7.     if (s->setup_len > sizeof(s->data_buf)) {
8.         fprintf(stderr,
9.             "usb_generic_handle_packet: ctrl buffer too
small (%d > %zu)\n",
10.            s->setup_len, sizeof(s->data_buf));
11.         p->status = USB_RET_STALL;
12.         return;
13.     }
14.     if (s->setup_buf[0] & USB_DIR_IN) {
15.         ...
16.     }else{
17.         s->setup_state = SETUP_STATE_DATA;
```



Get the length of setting data

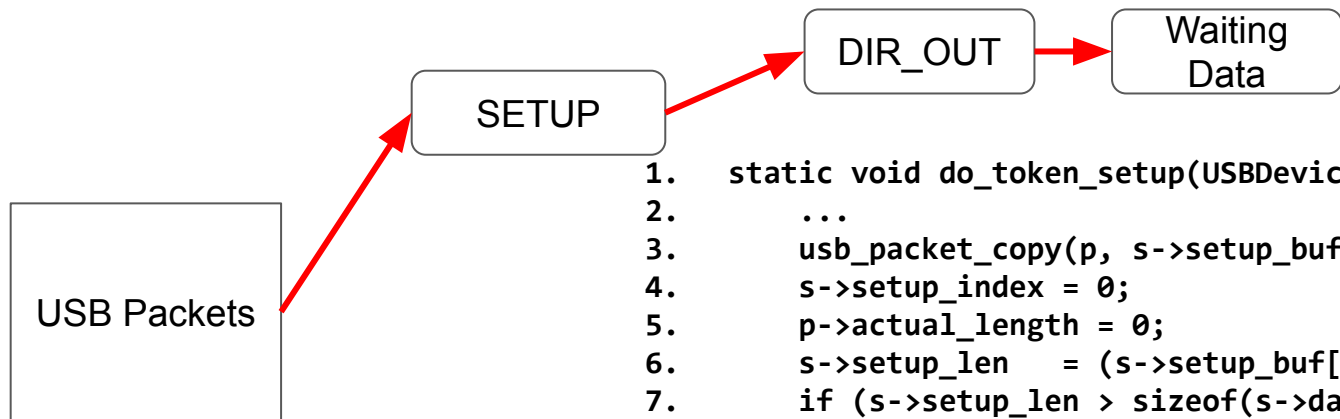
```

1. static void do_token_setup(USBDevice *s, USBPacket *p){
2.     ...
3.     usb_packet_copy(p, s->setup_buf, p->iov.size);
4.     s->setup_index = 0;
5.     p->actual_length = 0;
6.     s->setup_len = (s->setup_buf[7] << 8) | s->setup_buf[6];
7.     if (s->setup_len > sizeof(s->data_buf)) {
8.         fprintf(stderr,
9.             "usb_generic_handle_packet: ctrl buffer too
small (%d > %zu)\n",
10.             s->setup_len, sizeof(s->data_buf));
11.         p->status = USB_RET_STALL;
12.         return;
13.     }
14.     if (s->setup_buf[0] & USB_DIR_IN) {
15.         ...
16.     }else{
17.         s->setup_state = SETUP_STATE_DATA;
  
```



Check the input length if bigger than the buffer

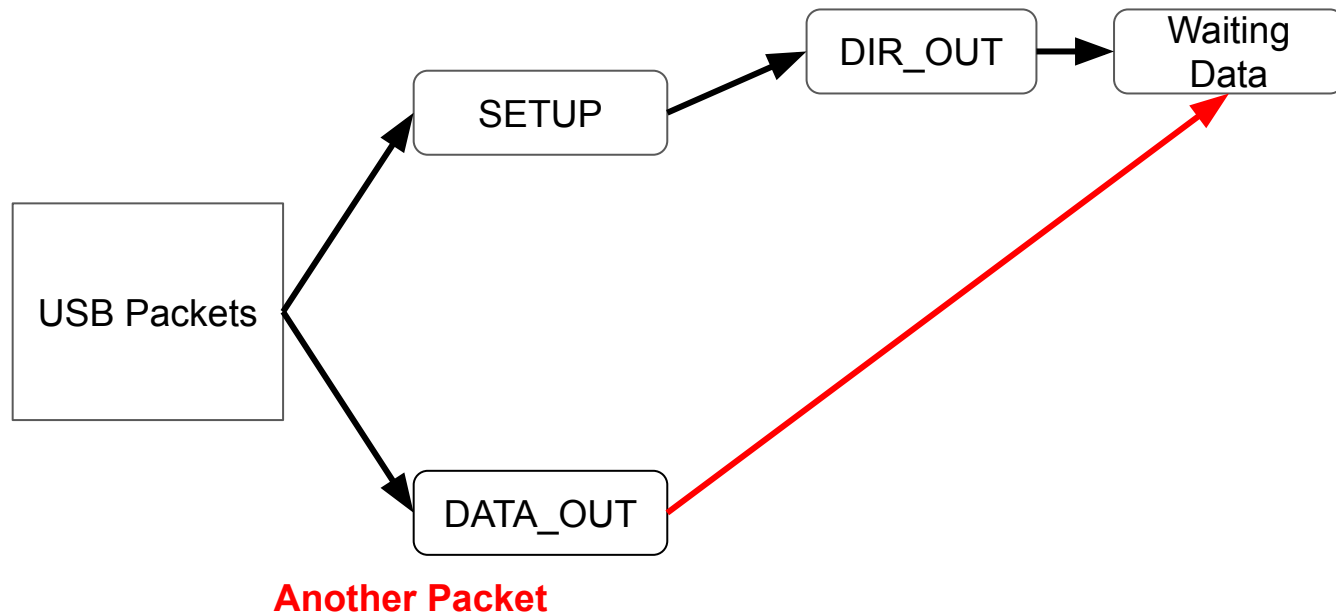
```
1. static void do_token_setup(USBDevice *s, USBPacket *p){
2.     ...
3.     usb_packet_copy(p, s->setup_buf, p->iov.size);
4.     s->setup_index = 0;
5.     p->actual_length = 0;
6.     s->setup_len = (s->setup_buf[7] << 8) | s->setup_buf[6];
7.     if (s->setup_len > sizeof(s->data_buf)) {
8.         fprintf(stderr,
9.             "usb_generic_handle_packet: ctrl buffer too
small (%d > %zu)\n",
10.             s->setup_len, sizeof(s->data_buf));
11.         p->status = USB_RET_STALL;
12.         return;
13.     }
14.     if (s->setup_buf[0] & USB_DIR_IN) {
15.         ...
16.     }else{
17.         s->setup_state = SETUP_STATE_DATA;
```

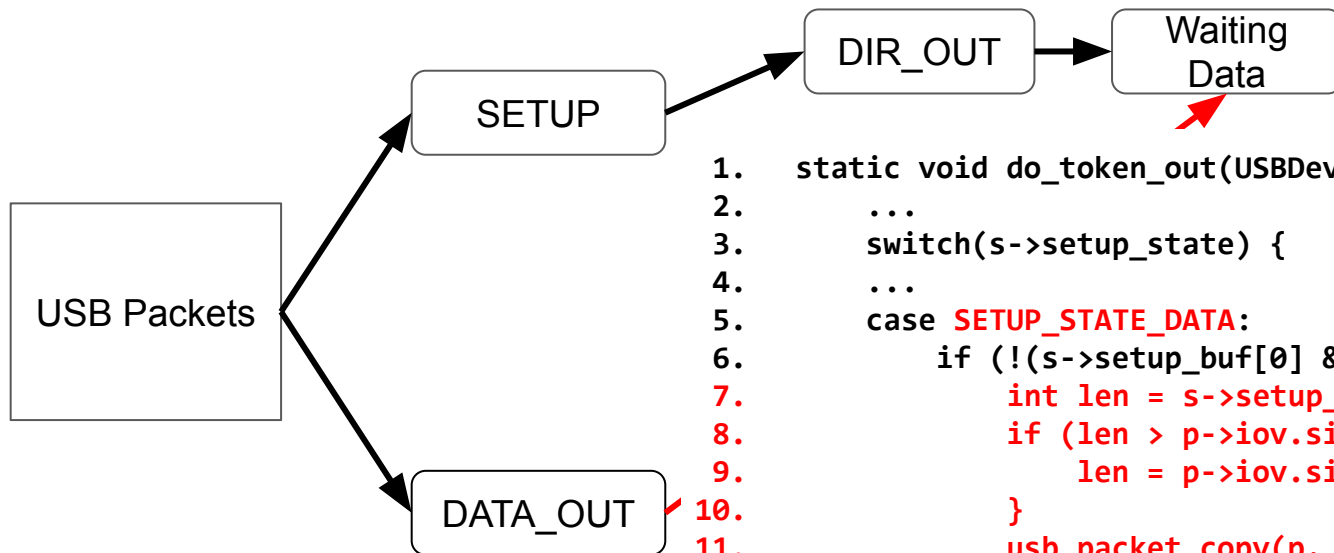


Return if check fails without clearing s->setup_len ! ! !

```

1. static void do_token_setup(USBDevice *s, USBPacket *p){
2.     ...
3.     usb_packet_copy(p, s->setup_buf, p->iov.size);
4.     s->setup_index = 0;
5.     p->actual_length = 0;
6.     s->setup_len = (s->setup_buf[7] << 8) | s->setup_buf[6];
7.     if (s->setup_len > sizeof(s->data_buf)) {
8.         fprintf(stderr,
9.             "usb_generic_handle_packet: ctrl buffer too
small (%d > %zu)\n",
10.            s->setup_len, sizeof(s->data_buf));
11.         p->status = USB_RET_STALL;
12.         return;
13.     }
14.     if (s->setup_buf[0] & USB_DIR_IN) {
15.         ...
16.     }else{
17.         s->setup_state = SETUP_STATE_DATA;
  
```



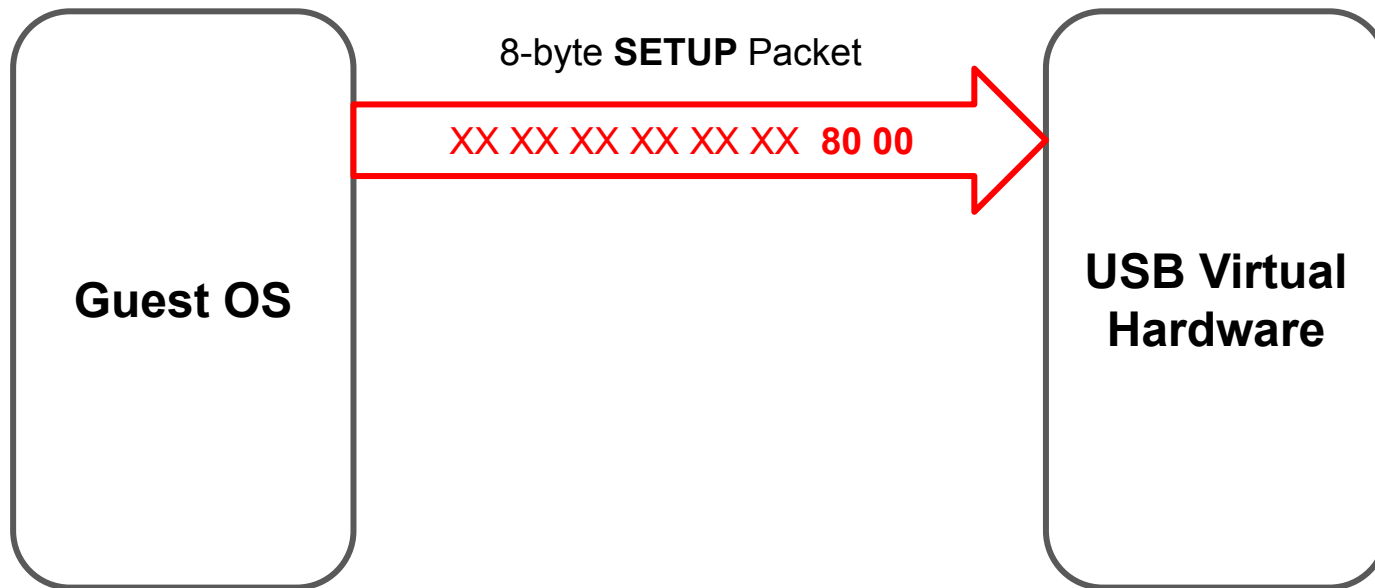


s->setup_len will be used in copying data

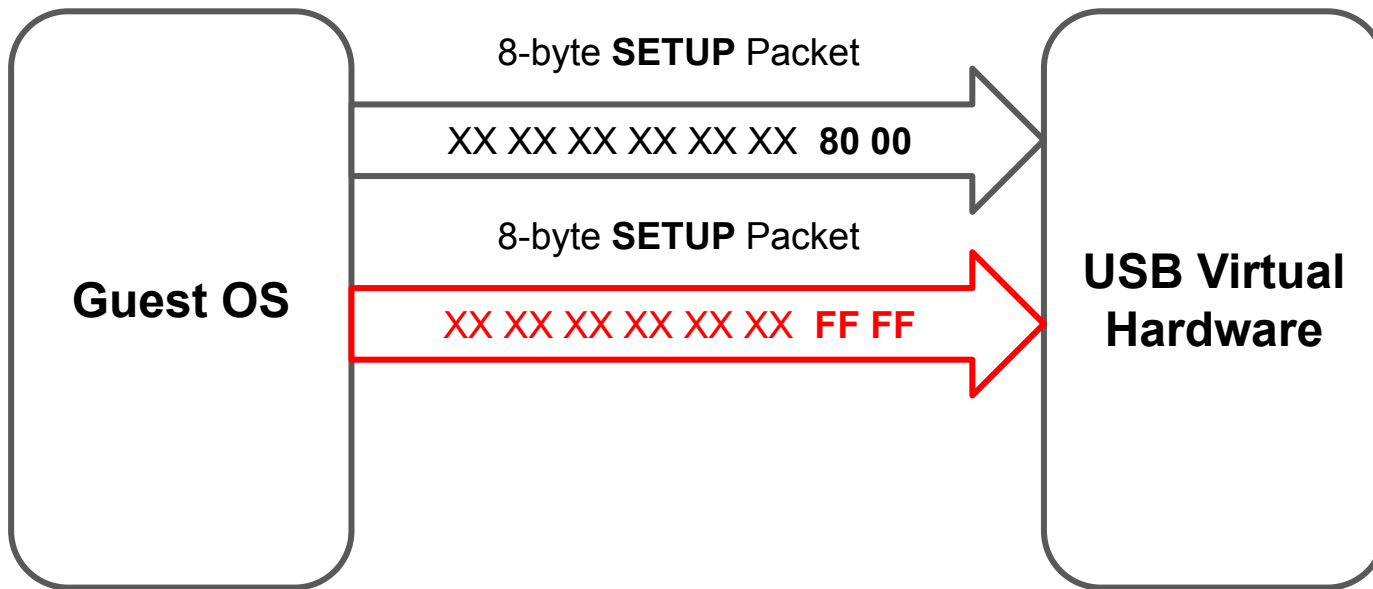
```

1. static void do_token_out(USBDevice *s, USBPacket *p){
2.     ...
3.     switch(s->setup_state) {
4.     ...
5.     case SETUP_STATE_DATA:
6.         if (!(s->setup_buf[0] & USB_DIR_IN)) {
7.             int len = s->setup_len - s->setup_index;
8.             if (len > p->iov.size) {
9.                 len = p->iov.size;
10.            }
11.            usb_packet_copy(p, s->data_buf + s->setup_index,
12.                len);
13.            s->setup_index += len;
14.            if (s->setup_index >= s->setup_len) {
15.                s->setup_state = SETUP_STATE_ACK;
16.            }
17.            return;
18.        }
    }

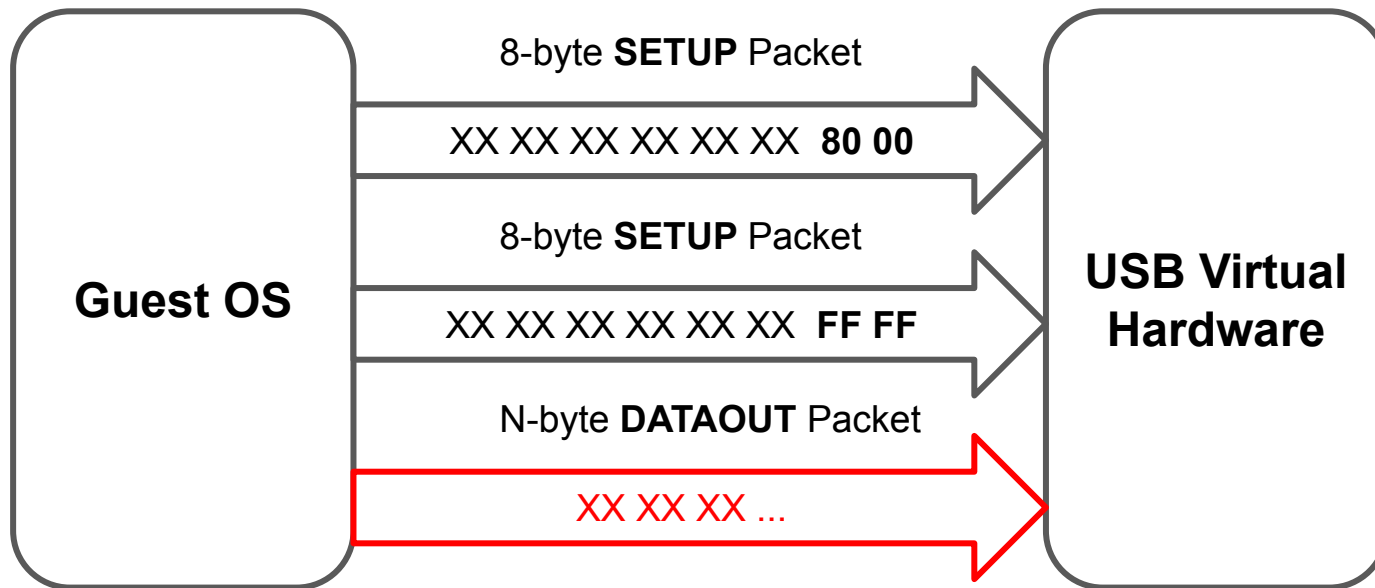
```

First, we send an 8-byte SETUP packet to make the `s->setup_state` be `SETUP_STATE_DATA`



Next, we send another 8-byte SETUP packet to make the `s->setup_len` big enough



Finally, we send another N-byte DATAOUT packet. We will overflow-write to the data buffer.



Black Box Escape

- It's hard for an attacker to get following information
 - Qemu's version
 - The binary file of Qemu

- We can do out-of-bound read and write of the databuf between 0-0xffff.

```
1. struct USBDevice {
2.     DeviceState qdev;
3.     ...
4.     uint8_t setup_buf[8];
5.     uint8_t data_buf[4096];
6.     int32_t remote_wakeup;
7.     int32_t setup_state;
8.     int32_t setup_len;
9.     int32_t setup_index;
10.
11.     USBEndpoint ep_ctl;
12.     USBEndpoint ep_in[USB_MAX_ENDPOINTS];
13.     USBEndpoint ep_out[USB_MAX_ENDPOINTS];
14.
15.     QLIST_HEAD(, USBDescString) strings;
16.     const USBDesc *usb_desc;
17.     /* Overrides class usb_desc if not NULL */
18.     ...
```

- We can do out-of-bound read and write of the databuf between 0-0xffff.
- How to leak some key information?

```
1. struct USBDevice {
2.     DeviceState qdev;
3.     ...
4.     uint8_t setup_buf[8];
5.     uint8_t data_buf[4096];
6.     int32_t remote_wakeup;
7.     int32_t setup_state;
8.     int32_t setup_len;
9.     int32_t setup_index;
10.
11.     USBEndpoint ep_ctl;
12.     USBEndpoint ep_in[USB_MAX_ENDPOINTS];
13.     USBEndpoint ep_out[USB_MAX_ENDPOINTS];
14.
15.     QLIST_HEAD(, USBDescString) strings;
16.     const USBDesc *usb_desc;
17.     /* Overrides class usb_desc if not NULL */
18.     ...
```

- We can do out-of-bound read and write of the databuf between 0-0xffff.
- How to leak some key information?
- The `usb_desc` contains the description of this USB device.

```
1. struct USBDevice {
2.     DeviceState qdev;
3.     ...
4.     uint8_t setup_buf[8];
5.     uint8_t data_buf[4096];
6.     int32_t remote_wakeup;
7.     int32_t setup_state;
8.     int32_t setup_len;
9.     int32_t setup_index;
10.
11.     USBEndpoint ep_ctl;
12.     USBEndpoint ep_in[USB_MAX_ENDPOINTS];
13.     USBEndpoint ep_out[USB_MAX_ENDPOINTS];
14.
15.     QLIST_HEAD(, USBDescString) strings;
16.     const USBDesc *usb_desc;
17.     /* Overrides class usb_desc if not NULL */
18.     ...
```


OOB Read and Write → Arbitrary Address Read CHAIR

- We can get the USBDescID by sending some USB packets.
- Arbitrary Address Read
 - Overwrite the pointer of USBDesc.
 - Get the USBDescID back.

```
1. struct USBDesc {
2.     USBDescID id;
3.     const USBDescDevice *full;
4.     const USBDescDevice *high;
5.     const USBDescDevice *super;
6.     const char* const *str;
7.     const USBDescMSOS *msos;
8. };
9.
10. struct USBDescID {
11.     uint16_t idVendor;
12.     uint16_t idProduct;
13.     uint16_t bcdDevice;
14.     uint8_t iManufacturer;
15.     uint8_t iProduct;
16.     uint8_t iSerialNumber;
17. };
```

- We get the address of USBDevice by reading the **USBEndpoint(ep_ctl, ep_in or ep_out)**

```
1. struct USBDevice {
2.     DeviceState qdev;
3.     ...
4.     uint8_t setup_buf[8];
5.     uint8_t data_buf[4096];
6.     int32_t remote_wakeup;
7.     int32_t setup_state;
8.     int32_t setup_len;
9.     int32_t setup_index;
10.
11.     USBEndpoint ep_ctl;
12.     USBEndpoint ep_in[USB_MAX_ENDPOINTS];
13.     USBEndpoint ep_out[USB_MAX_ENDPOINTS];
14.
15.     QLIST_HEAD(, USBDescString) strings;
16.     const USBDesc *usb_desc;
17.     /* Overrides class usb_desc if not NULL */
18.     ...
19.
```

- We get the address of USBDevice by reading the **USBEndpoint(ep_ctl, ep_in or ep_out)**

```
1. struct USBEndpoint {
2.     uint8_t nr;
3.     uint8_t pid;
4.     uint8_t type;
5.     uint8_t ifnum;
6.     int max_packet_size;
7.     int max_streams;
8.     bool pipeline;
9.     bool halted;
10.    USBDevice *dev;
11.    QTAILQ_HEAD(, USBPacket) queue;
12. };
```

- We get the address of USBDevice by reading the **USBEndpoint(ep_ctl, ep_in or ep_out)**
- **DeviceState** has a free function pointers

```
1. struct USBDevice {
2.     DeviceState qdev;
3.     ...
4.     uint8_t setup_buf[8];
5.     uint8_t data_buf[4096];
6.     int32_t remote_wakeup;
7.     int32_t setup_state;
8.     int32_t setup_len;
9.     int32_t setup_index;
10.
11.     USBEndpoint ep_ctl;
12.     USBEndpoint ep_in[USB_MAX_ENDPOINTS];
13.     USBEndpoint ep_out[USB_MAX_ENDPOINTS];
14.
15.     QLIST_HEAD(, USBDescString) strings;
16.     const USBDesc *usb_desc;
17.     /* Overrides class usb_desc if not NULL */
18.     ...
19. }
```

- We get the address of USBDevice by reading the **USBEndpoint(ep_ctl, ep_in or ep_out)**
- **DeviceState** has a free function pointers

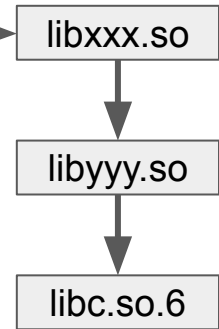
```
1. struct DeviceState {
2.     /*< private >*/
3.     Object parent_obj;
4.     ...
5. };
6.
7. struct Object
8. {
9.     /*< private >*/
10.    ObjectClass *class;
11.    ObjectFree *free;
12.    GHashTable *properties;
13.    uint32_t ref;
14.    Object *parent;
15. };
16.
```

Arbitrary Address Read → Leak libc

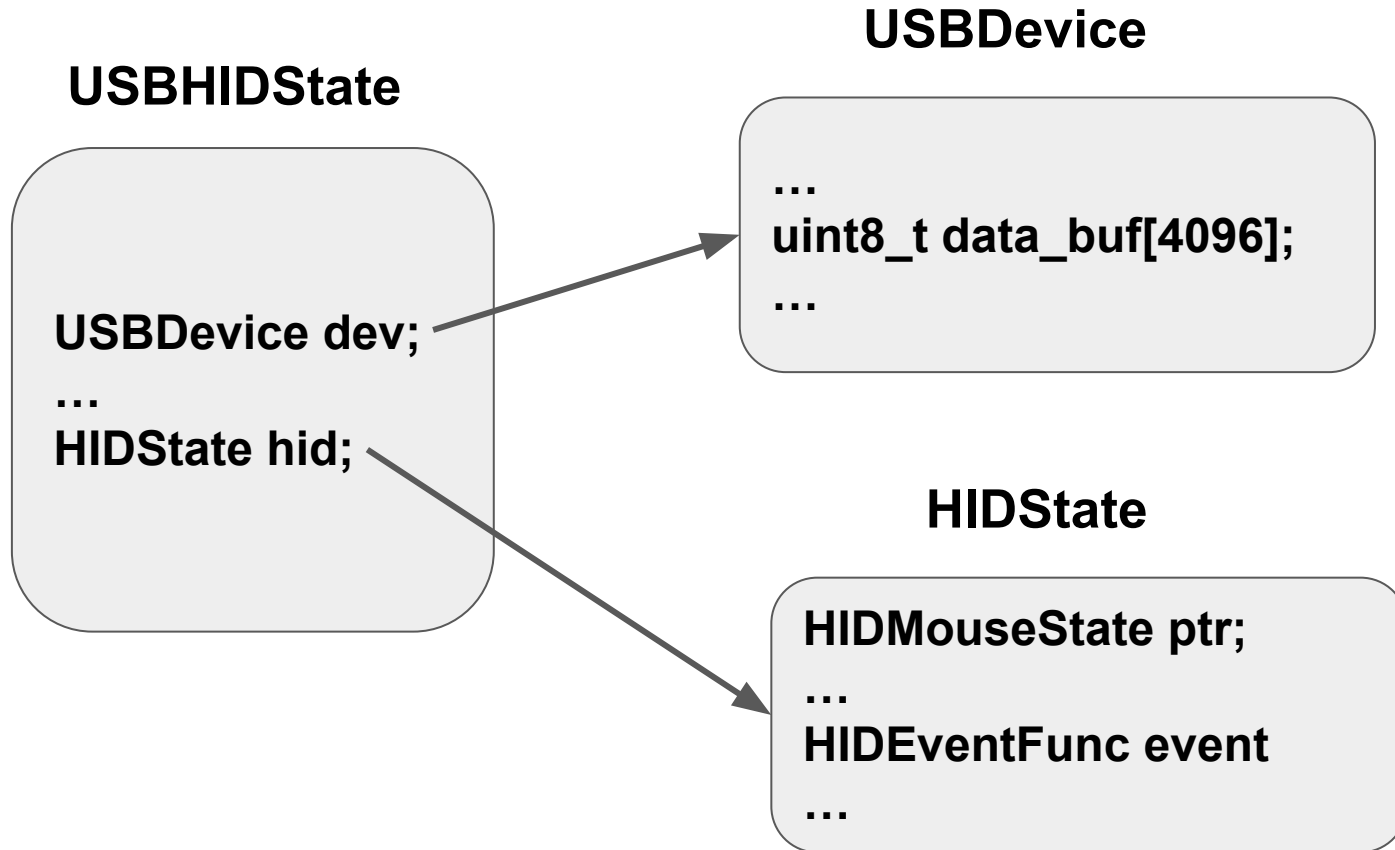


- We get the address of USBDevice by reading the **USBEndpoint(ep_ctl, ep_in or ep_out)**
- **DeviceState** has a free function pointers
- We will finally get the free address in libc

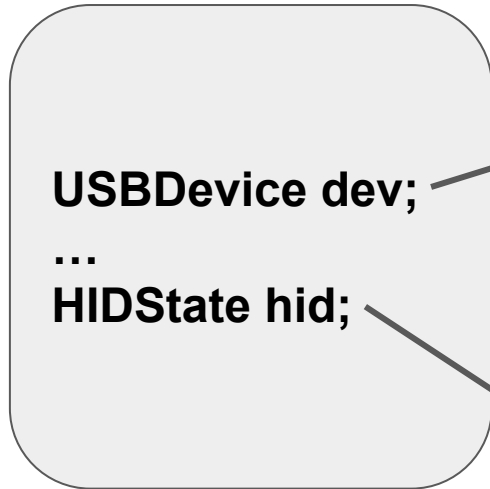
```
1. struct DeviceState {
2.     /*< private >*/
3.     Object parent_obj;
4.     ...
5. };
6.
7. struct Object
8. {
9.     /*< private >*/
10.    ObjectClass *class;
11.    ObjectFree *free;
12.    GHashTable *properties;
13.    uint32_t ref;
14.    Object *parent;
15. };
16.
```



- After getting address of free, we get the address of system like pwntools.Dynelf does.
 - Search ELF magic number forward first to get the base address of libc
 - Find **.dynstr** and **.dynsym** section
 - Find “**system**” in **.dynstr** and get the offset in **.dynsym**

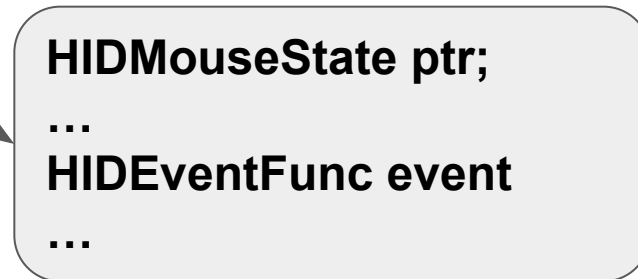


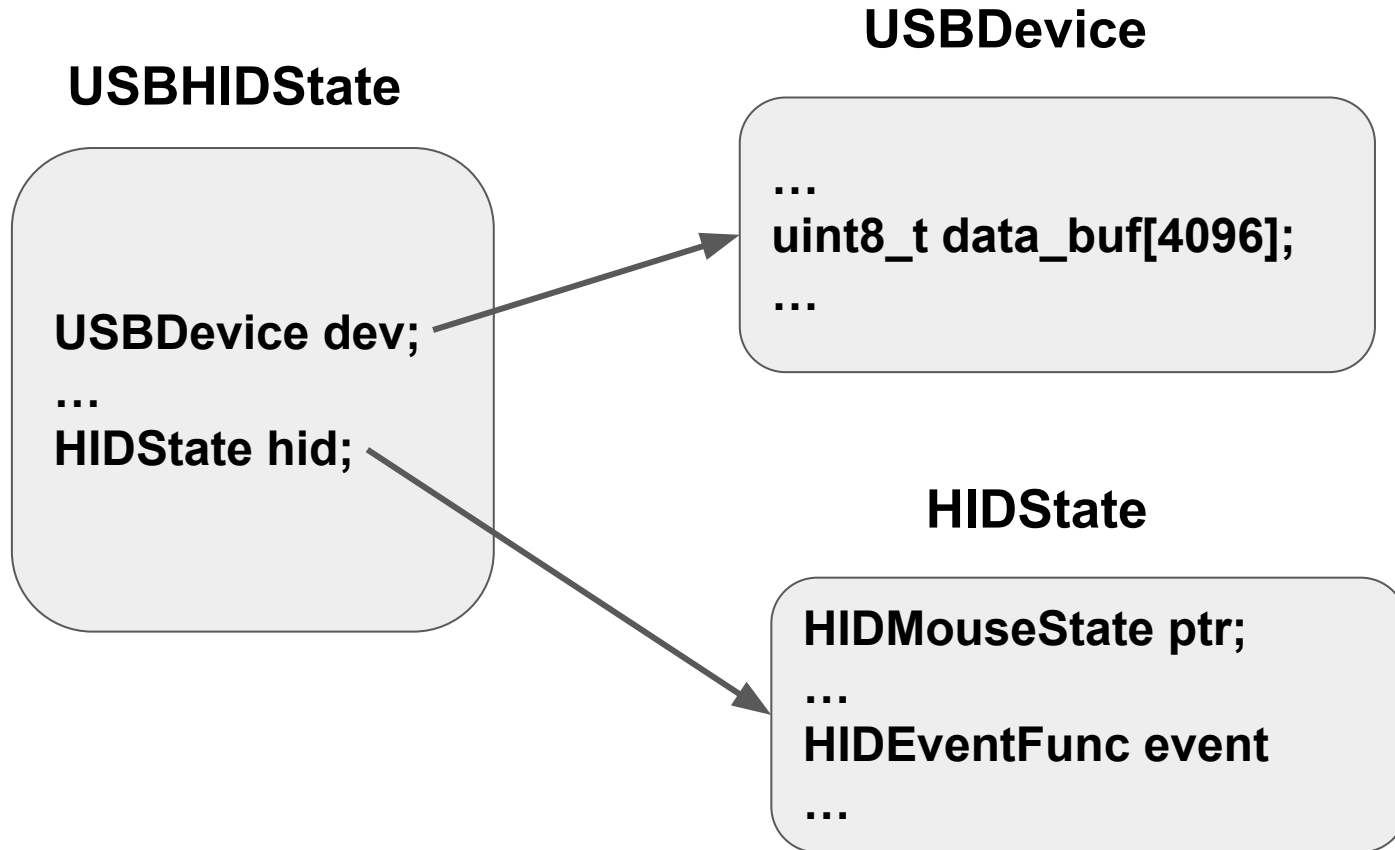
USBHIDState



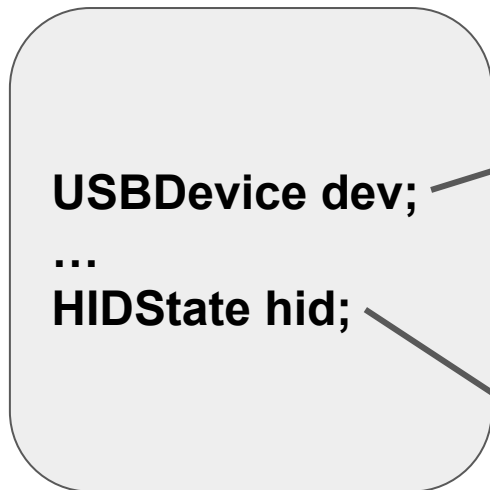
```
1. static void hid_idle_timer(void *opaque)
2. {
3.     HIDState *hs = opaque;
4.
5.     hs->idle_pending = true;
6.     hs->event(hs);
7. }
```

HIDState

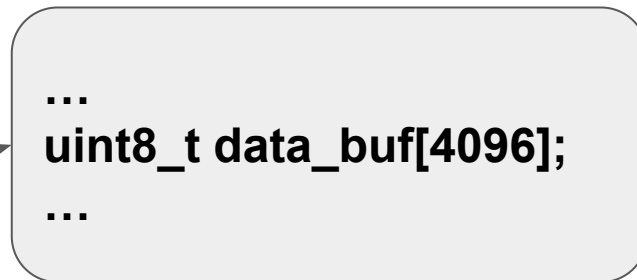




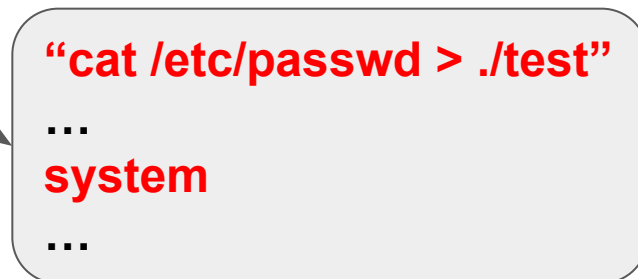
USBHIDState



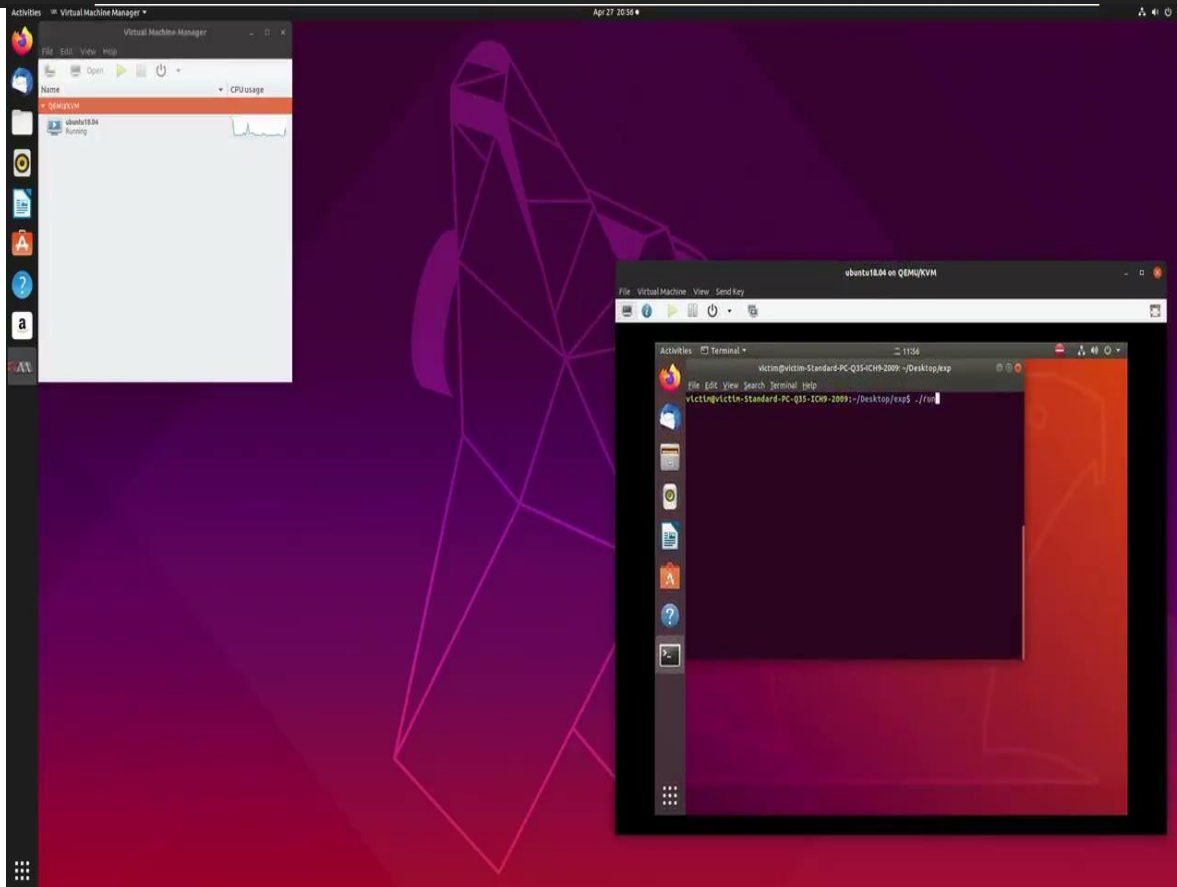
USBDevice



HIDState



Demo



- Sandboxes are necessary even in public cloud environments
- Good vulnerabilities can do a lot of interesting things
- The skills used in CTF are helpful

Thanks!



@flyYY_



长亭科技
CHAITIN