

GEEKCON

「 30+5 深度分享

30 + 5 In-depth Sharing

Xuan Xing & Eugene Rodionov

」

# A conversation happened long time ago...

- My manager:
  - Maybe we should fuzz the Linux HID driver?
  - Any idea how to do it?
- <if I were an experienced kernel hacker>:
  - syzkaller?
- Actual me:
  - `clang -fsanitize=fuzzer drivers/hid/hid-core.c kernel_stubs.c .....`
- Clang:
  - fatal error: too many errors emitted, stopping now [-ferror-limit=]
- Matt Alexander (former Googler/RedTeamer):
  - Oh, you want to run kernel code in userspace? try LKL!

# Bypassing Kernel Barriers

Fuzzing Linux Kernel in Userspace with LKL

Xuan Xing & Eugene Rodionov

24-October-2024

Google

 Agenda

- Introduction
- LKL Basics
- Fuzzing on LKL
- LKL Fuzzers
- Quick Demo
- Conclusion

GEEKCON

# Introduction

# Who are we?

The Android Red Team is tasked to increase **Android** security by attacking key Android components and features, identifying critical vulnerabilities before adversaries.

Our work includes:

- Offensive Security Reviews to verify (break) security assumptions
- Scale through tool development (e.g. continuous fuzzing)
- Develop proof of concepts to demonstrate real-world impact
- Assess the efficacy of security mitigations



GEEKCON

What is LKL?

# LKL Overview

## Linux Kernel Library (LKL)<sup>1</sup> builds Linux kernel as a user-space library

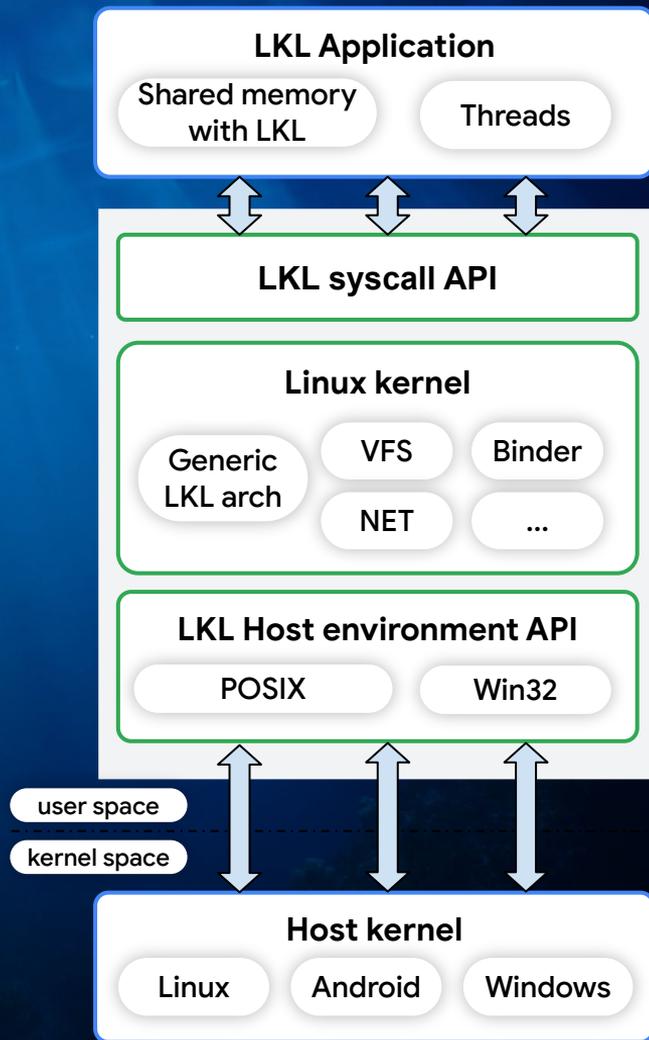
- Implemented as Linux arch-port
- LKL vs UML

## LKL building blocks

- Host environment API -- portability layer
- Linux kernel code
- LKL syscall API exposed to the user-space application

## Run kernel code without launching a VM

- Kernel unit testing
- Fuzzing!<sup>2,3</sup>



[1] <https://github.com/lkl/linux>

[2] Xu et al., Fuzzing File Systems via Two-Dimensional Input Space Exploration

[3] <https://github.com/atrosinenko/kbdysch>

# Using LKL from your C program

```
int ret = lk1_start_kernel(&lk1_host_ops, "mem=50M");

lk1_mount_fs("sysfs");
lk1_mount_fs("proc");
lk1_mount_fs("dev");

int binder_fd = lk1_sys_open("/dev/binder", O_RDWR | O_CLOEXEC, 0);
void *binder_map = lk1_sys_mmap(NULL, BINDER_VM_SIZE,
                                PROT_READ, MAP_PRIVATE, binder_fd, 0);

struct lk1_binder_version version = { 0 };
ret = lk1_sys_ioctl(binder_fd, LKL_BINDER_VERSION, &version);
```

# Using LKL from your C program

```
int ret = lkl_start_kernel(&lkl_host_ops, "mem=50M");

lkl_mount_fs("sysfs");
lkl_mount_fs("proc");
lkl_mount_fs("dev");

int binder_fd = lkl_sys_open("/dev/binder", O_RDWR | O_CLOEXEC, 0);
void *binder_map = lkl_sys_mmap(NULL, BINDER_VM_SIZE,
                                PROT_READ, MAP_PRIVATE, binder_fd, 0);

struct lkl_binder_version version = { 0 };
ret = lkl_sys_ioctl(binder_fd, LKL_BINDER_VERSION, &version);
```

# Using LKL from your C program

```
int ret = lkl_start_kernel(&lkl_host_ops, "mem=50M");

lkl_mount_fs("sysfs");
lkl_mount_fs("proc");
lkl_mount_fs("dev");

int binder_fd = lkl_sys_open("/dev/binder", O_RDWR | O_CLOEXEC, 0);
void *binder_map = lkl_sys_mmap(NULL, BINDER_VM_SIZE,
                                PROT_READ, MAP_PRIVATE, binder_fd, 0);

struct lkl_binder_version version = { 0 };
ret = lkl_sys_ioctl(binder_fd, LKL_BINDER_VERSION, &version);
```

# Using LKL from your C program

```
int ret = lkl_start_kernel(&lkl_host_ops, "mem=50M");

lkl_mount_fs("sysfs");
lkl_mount_fs("proc");
lkl_mount_fs("dev");

int binder_fd = lkl_sys_open("/dev/binder", O_RDWR | O_CLOEXEC, 0);
void *binder_map = lkl_sys_mmap(NULL, BINDER_VM_SIZE,
                                PROT_READ, MAP_PRIVATE, binder_fd, 0);

struct lkl_binder_version version = { 0 };
ret = lkl_sys_ioctl(binder_fd, LKL_BINDER_VERSION, &version);
```

# Applications

- Built-in applications
  - tests/boot: testing LKL booting
  - fs2tar: converting file system image to a tar archive
  - cptofs/cpfromfs: copies files to/from a file system image
- Kernel fuzzing
  - kBdysch: A collection of kernel fuzzers in userspace
  - Janus: Fuzzing file system using lkl
  - HID fuzzer: A sample HID fuzzer developed by Android Red Team
- Other interesting ideas
  - LKL.js: running Linux kernel on JavaScript directly
  - User Space TCP: a full TCP stack in userspace for better security

GEEKCON

# Enabling Fuzzing on LKL

# Anatomy of LKL fuzzer

## LKL enables fuzzing Linux kernel code in user-space

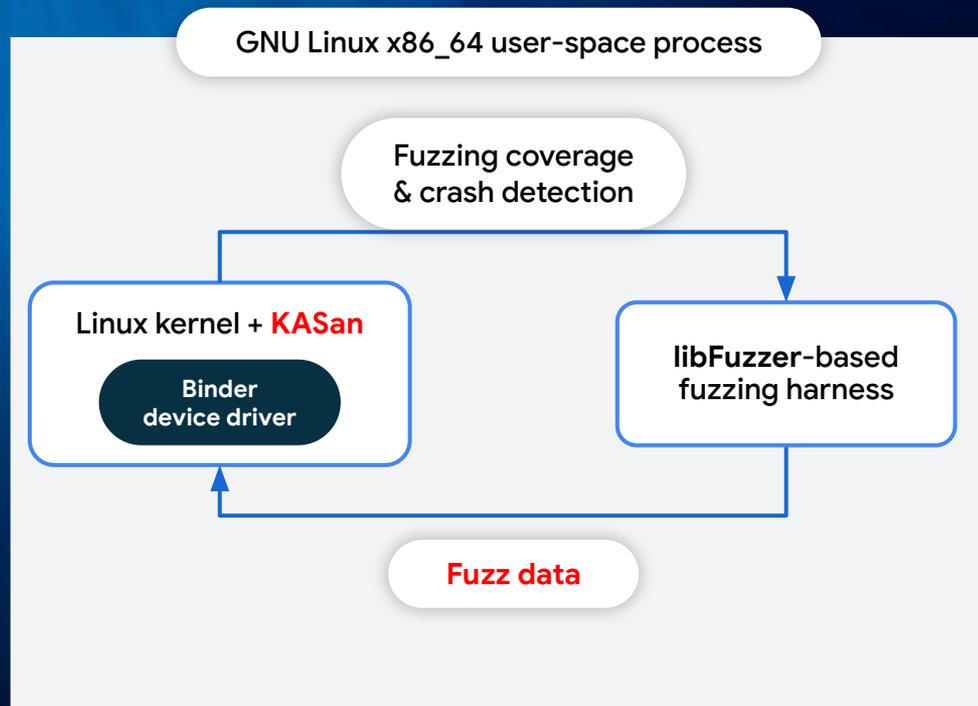
- Use in-process fuzzing engine, such as **libFuzzer**

## Advantages

- High fuzzing performance on x86\_64
- Ease of custom modifications
  - e.g. mocking hardware, custom scheduler(?)

## Limitations

- No SMP in LKL
- x86\_64 vs aarch64 -- potential false positives, false negatives

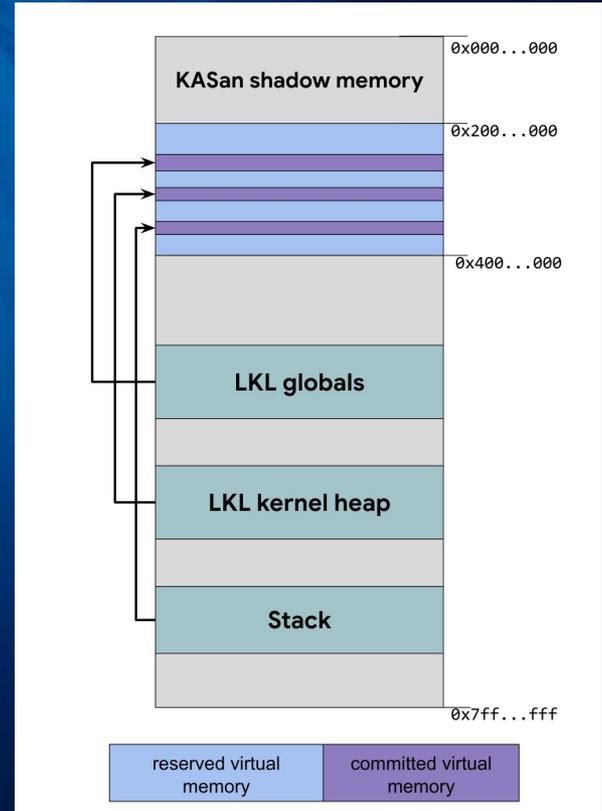


# Support added to LKL for fuzzing

- **LibFuzzer support**
  - Coverage guided fuzzing
  - Source code coverage for report visualization
- **Address Sanitizer**
  - Catch memory issues during fuzzing
  - Study lead to a [Red Team blogpost](#) on bare metal kASAN implementation
- **Stacktrace and symbolization**
- **virtio support**
  - Allow fuzzing hardware drivers with virtio backend
- **Generic fuzzing related scripts**
  - Building fuzzers
  - Makefiles, etc

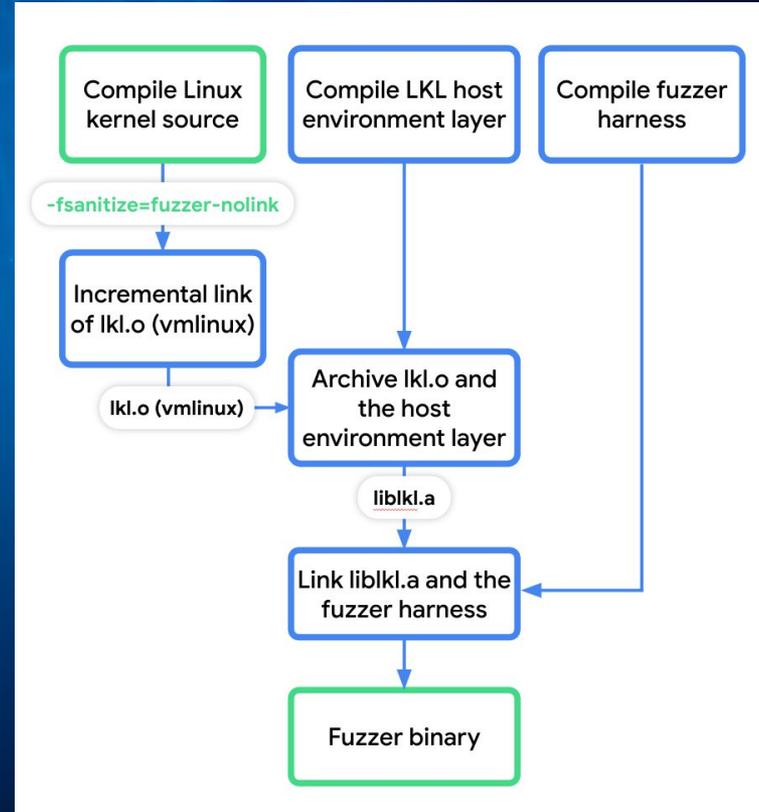
# LKL KASan details

- KASan provides actionable reports for invalid memory access:
  - OOB, user-after-free, double-free
  - covers stack, heap and globals
- User-space ASan in LKL:
  - ASan shadow memory poisoning routines are invoked in global constructors
  - Which might be problematic due to specifics of globals initialization in Linux kernel
- LKL implements generic KASan:
  - `-fsanitize=kernel-address`
  - arch-specific KASan implementation



# LKL Fuzzing Coverage

- LKL relies on libFuzzer-based fuzzing code coverage instrumentation
- KCOV is an alternative solution:
  - Needs additional implementation to feed the coverage feedback to libFuzzer engine



# LKL fuzzing limitations

- **Lack of MMU support**
  - Certain drivers can not be fuzzed
  - Fixable but requires resources
- **No SMP support**
  - Not easy to test certain concurrency scenarios
- **Challenges with coverage build**
  - Large binary size
  - Taking too long to build
- **Only support x86-64 arch**
  - Limited AArch/AArch64 support
  - Drivers specific to non-x86-64 architecture can't be fuzzed
- **Need to mock lower level interfaces**
  - Mocking hardware support
  - Using virtio backend
- **Not in linux main branch (yet)**
  - Requiring effort to be upstreamed

GEEKCON

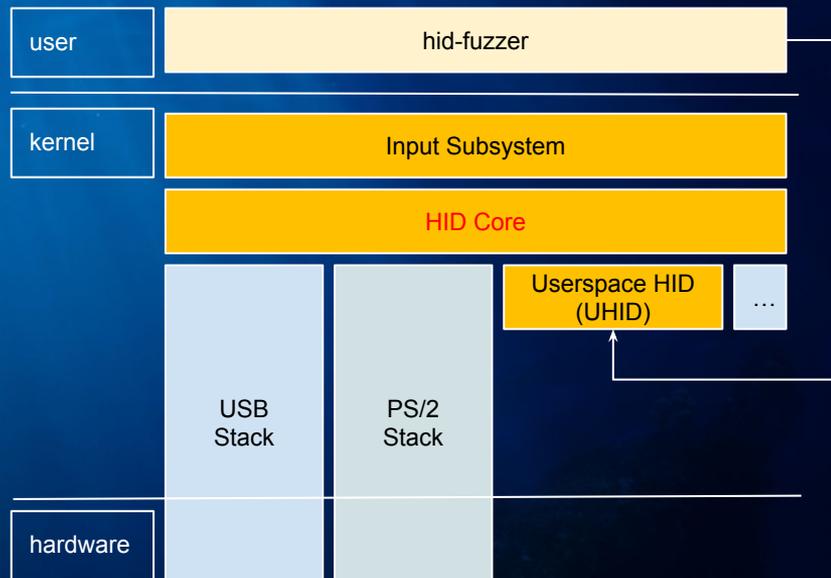
# LKL Fuzzers

# LKL fuzzers

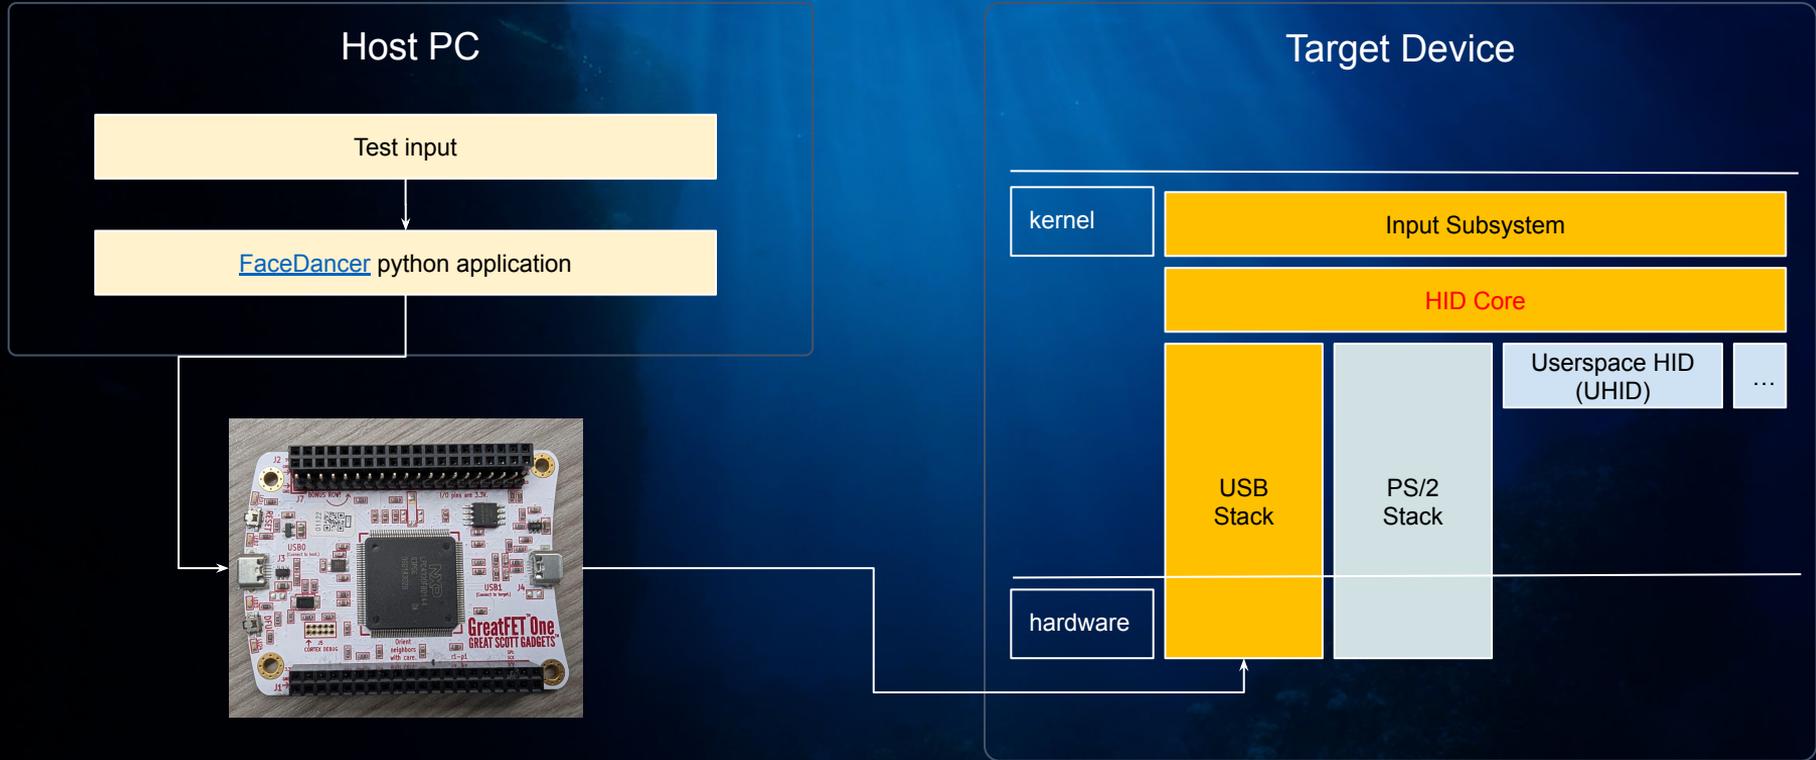
- Human Interface Device (HID) fuzzer
  - A sample LKL fuzzer targeting HID driver
- Binder fuzzer
  - Targeting Android Binder driver
- Virtio driver fuzzing
  - Used by Android pKVM with special attacking scenario (malicious host against guest VM)
  - Virtio based device drivers

# Human Interface Device (HID) fuzzing

- Target:
  - Linux HID system
  - Complex spec ([hid1\\_11.pdf](#), [hut1\\_5.pdf](#))
- Attacking scenario:
  - Physical access ([juice jacking](#))
- Details
  - Inspired by [kBdysch](#) fuzzing project
  - A proof of concept for LKL fuzzing support development
- Issues discovered and verified with USB simulation
  - Two bulletin class issues (CVE-2020-0465)
  - 6 stability issues



# Human Interface Device (HID) fuzzing (cont.)



# Exploitability of HID bugs

- We didn't attempt to exploit identified issues in HID driver
- However, ...



The image shows a slide from a Black Hat Europe 2021 briefing. The slide is split into two main sections: a white left side and a green right side. On the white side, the Riscure logo is displayed with the tagline 'driving your security forward'. On the green side, the Black Hat Europe 2021 logo is at the top, followed by the dates 'november 10-11, 2021' and the word 'BRIEFINGS'. The main title of the briefing is 'Achieving Linux Kernel Code Execution Through A Malicious USB Device'. Below the title, the names and titles of the presenters, Martijn Bogaard and Dana Geist, are listed along with their contact information.

**riscure**  
driving your security forward

**black hat**  
EUROPE 2021  
november 10-11, 2021  
BRIEFINGS

Achieving Linux Kernel Code Execution Through A Malicious USB Device

Martijn Bogaard  
Principal Security Analyst  
[martijn@riscure.com](mailto:martijn@riscure.com)  
@jmartijnb

Dana Geist  
Senior Security Analyst  
@geistdana

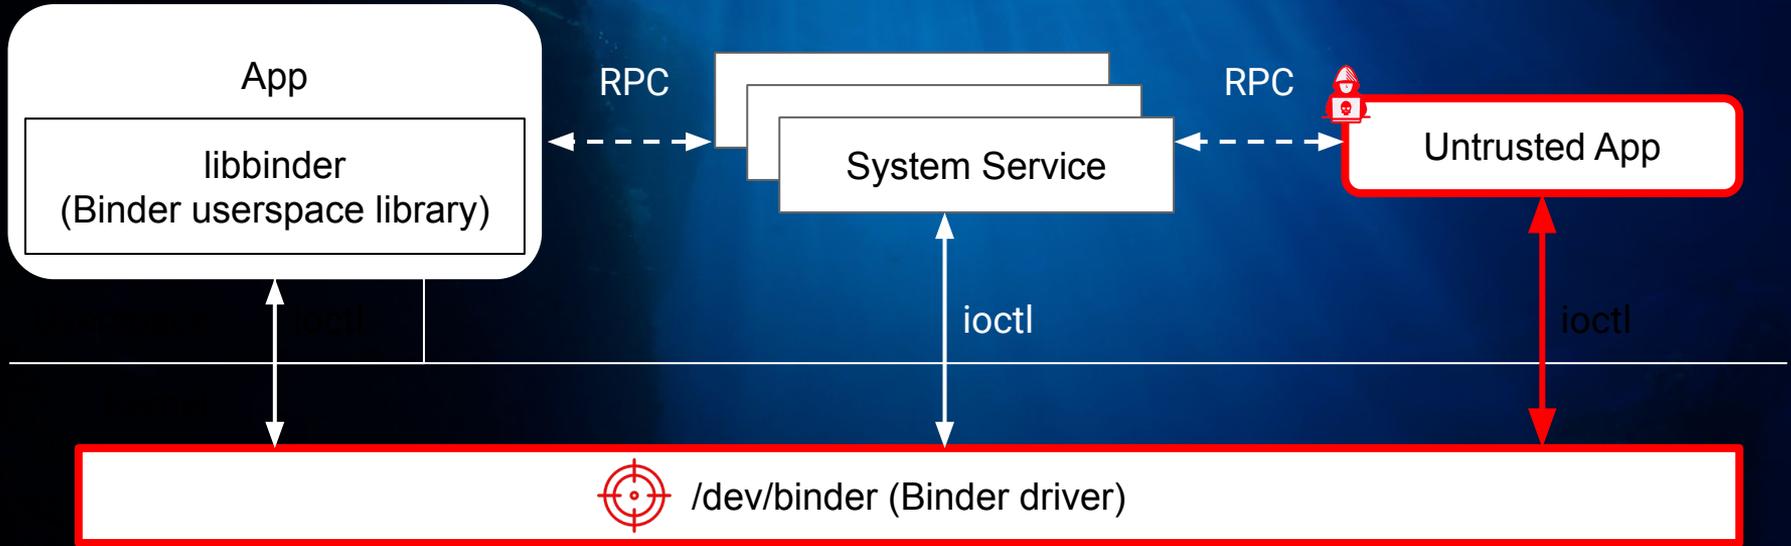
GEEKCON

# Android Binder Driver Fuzzing

# What is Binder?

- Primary inter-process communication (IPC) channel on Android
- Support passing file descriptors, objects containing pointers, etc.
- Composed of a userspace library (libbinder) and a kernel driver (/dev/binder)
- Provide Remote Procedure Call (RPC) framework for Java and C/C++

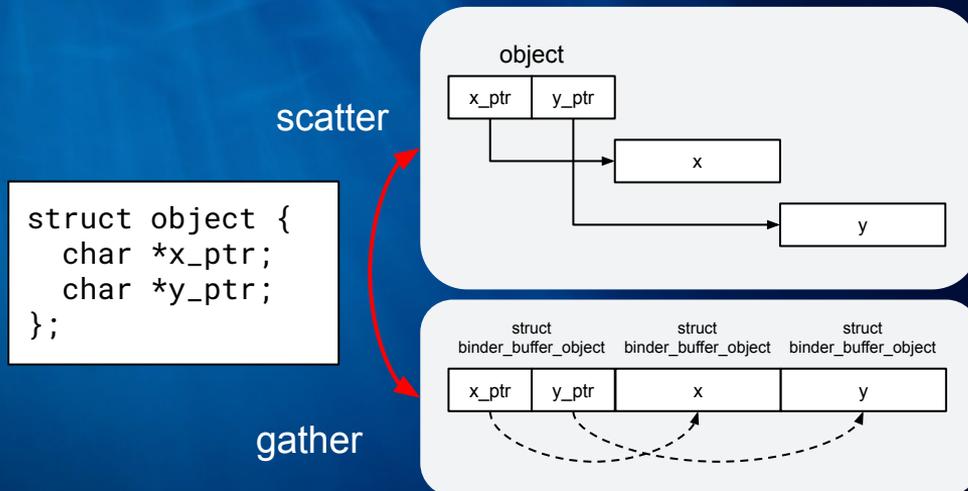
# Binder Threat Model



# Binder Fuzzing Challenges

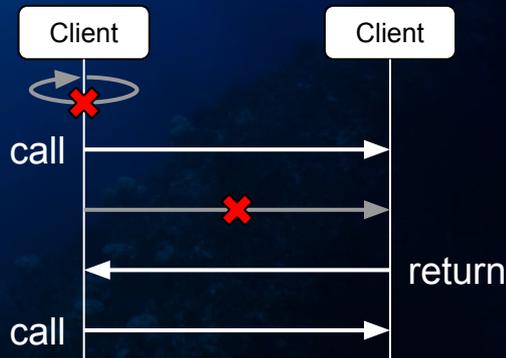
## Data dependencies

- Binder commands
- Scatter-gather data structures (BINDER\_TYPE\_PTR)



## State dependencies

- Synchronous IPC
  - Cannot send transaction to oneself
  - Multiple outstanding transactions not allowed

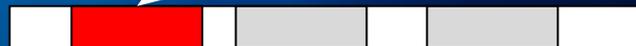


# Binder Fuzzing Challenges

## State dependencies

- Some inputs depend on previous IOCTL calls
  - e.g. Transaction buffers (BC\_FREE\_BUFFER)

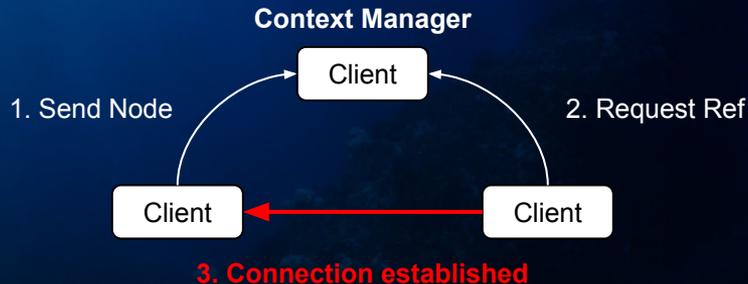
```
ioctl(binder_fd, BINDER_WRITE_READ, x) // 1.  
  
// y = x->read_buffer->...->buffer  
ioctl(binder_fd, BC_FREE_BUFFER, y) // 2.
```



Binder memory map

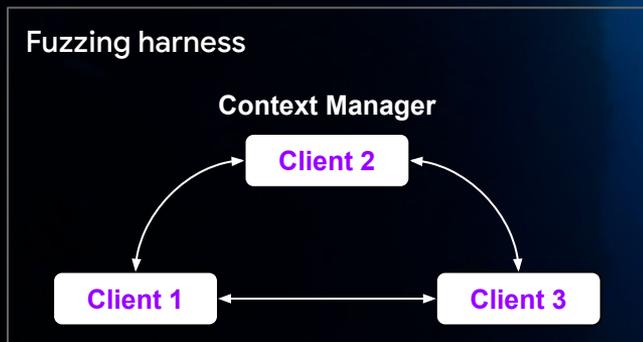
## Multi-process coordination

- All communication requires a **Context Manager**
- Node & Ref setup is required to establish a connection



# Fuzzing Harness

- Simulate IPC interactions between multiple clients
- 3 **clients** (1 Context Manager)
  - **IOCTL calls** and **data**



## Fuzz data

```
client_1 {
  binder_write {
    binder_commands {
      transaction {
        binder_objects { binder { ptr: 0xbeef } }
      }
    }
  }
}
client_2 {
  binder_read { ... }
  binder_write {
    binder_commands { free_buffer { ... } }
  }
}
client_3 { ... }
client_2 { ... }
client_3 { ... }
client_1 { ... }
```

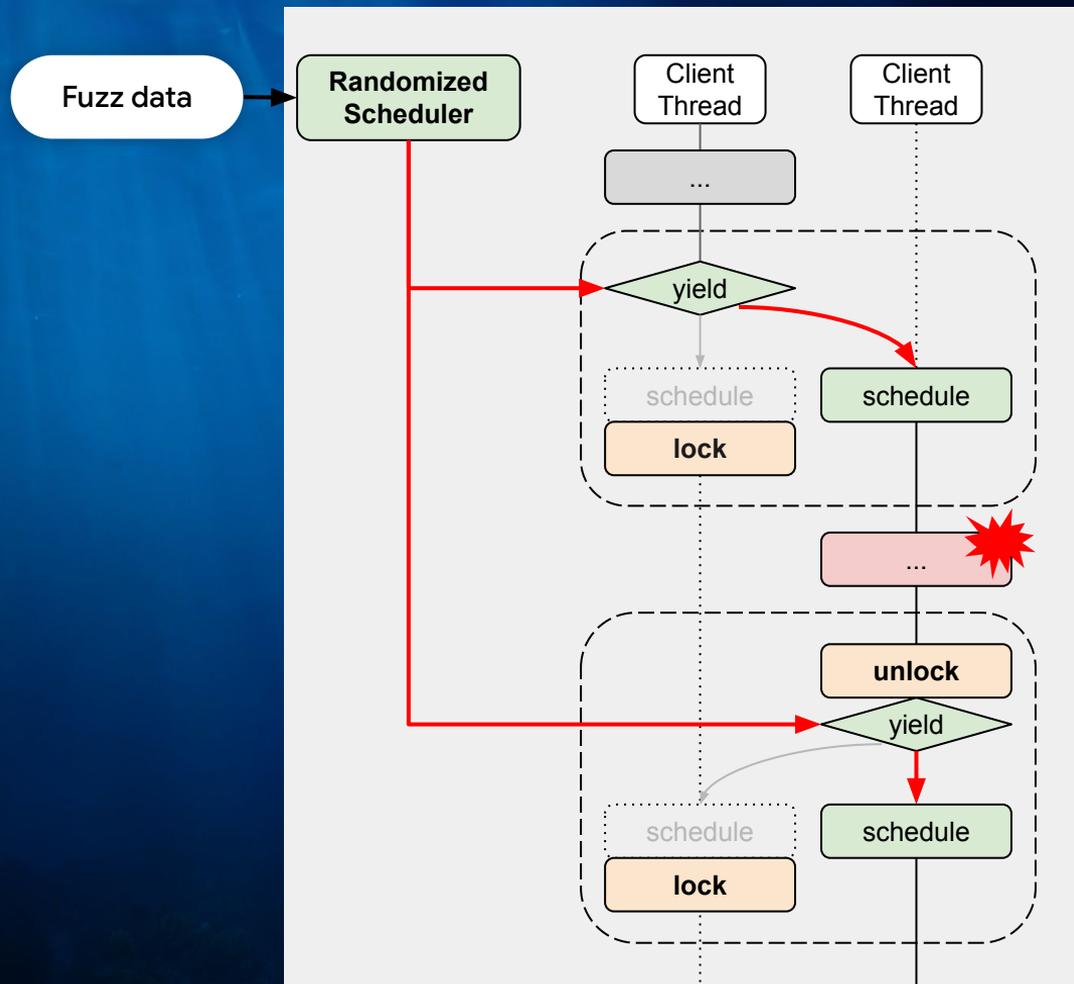
# Randomized Scheduler

**Deterministically** simulate thread interleaving based on fuzz data<sup>1</sup>

Insert yield points before/after synchronization primitives

- spin\_lock, spin\_unlock
- mutex\_lock, mutex\_unlock

[1] Williamson, N., *Catch Me If You Can: Deterministic Discovery of Race Conditions with Fuzzing*. Black Hat USA, (2022).



# Results

- Achieved 68% line coverage

## Coverage Report

Created: 2024-05-06 09:49

Click [here](#) for information about interpreting this report.

Filename	Function Coverage	Line Coverage	Region Coverage	Branch Coverage
<a href="#">binder.c</a>	79.14% (110/139)	68.55% (3086/4502)	50.90% (4486/8813)	50.50% (1317/2608)
<a href="#">binder_alloc.c</a>	78.38% (29/37)	70.41% (564/801)	49.87% (752/1508)	49.20% (185/376)
<a href="#">binder_alloc.h</a>	50.00% (1/2)	11.11% (1/9)	50.00% (1/2)	- (0/0)
<a href="#">binder_internal.h</a>	60.00% (3/5)	68.75% (11/16)	73.68% (14/19)	- (0/0)
<b>Totals</b>	<b>78.14% (143/183)</b>	<b>68.73% (3662/5328)</b>	<b>50.79% (5253/10342)</b>	<b>50.34% (1502/2984)</b>

Generated by llvm-cov -- llvm version 12.0.6git

- Discovered CVE-2023-20938 & CVE-2023-21255

30+5



*How to Fuzz Your Way to Android  
Universal Root: Attacking Android Binder*

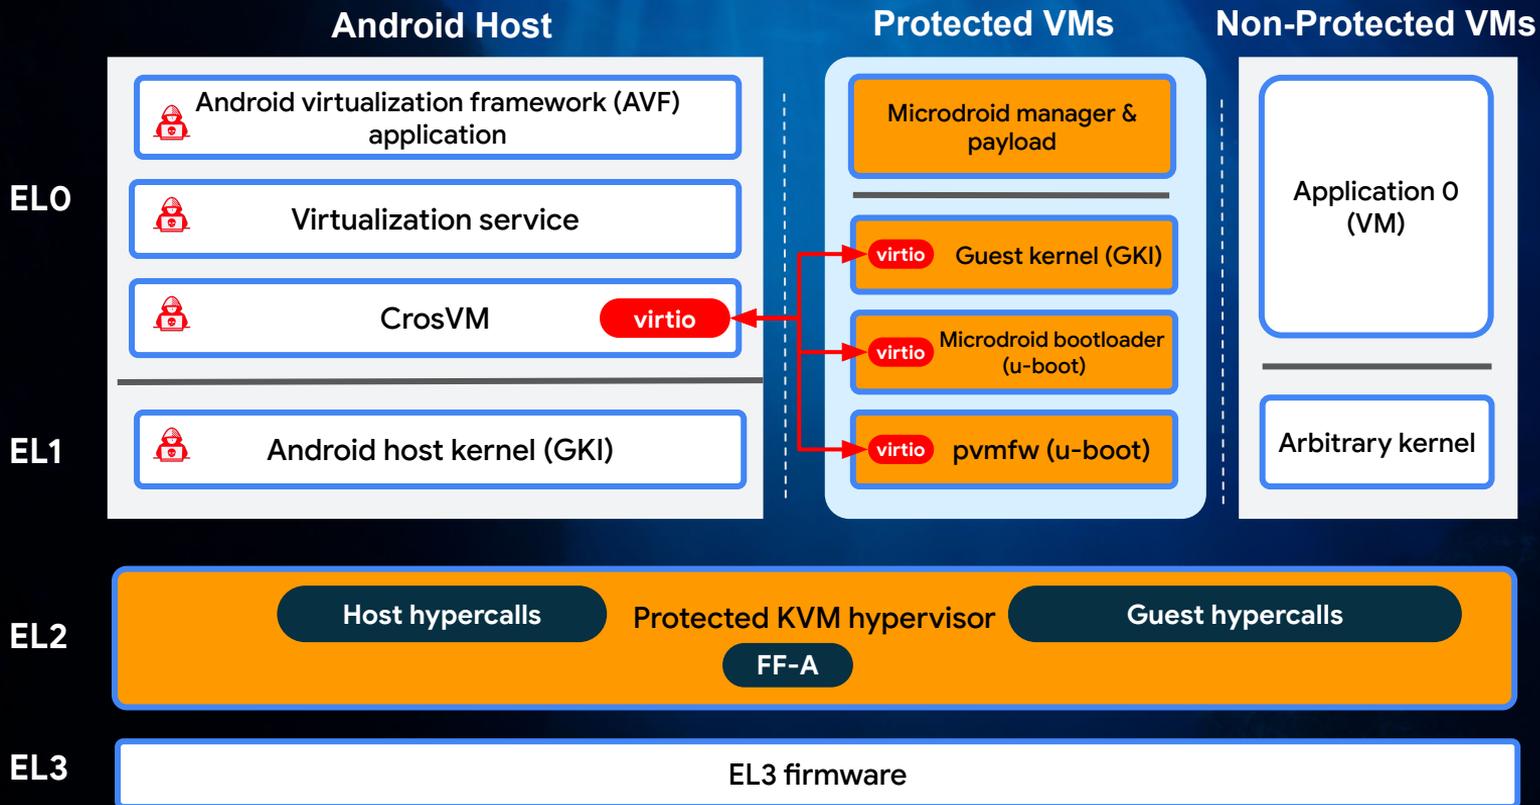
我是如何发现并利用底层通信漏洞控制整个安卓手机的

分享者: Google Android Red Team [美国]

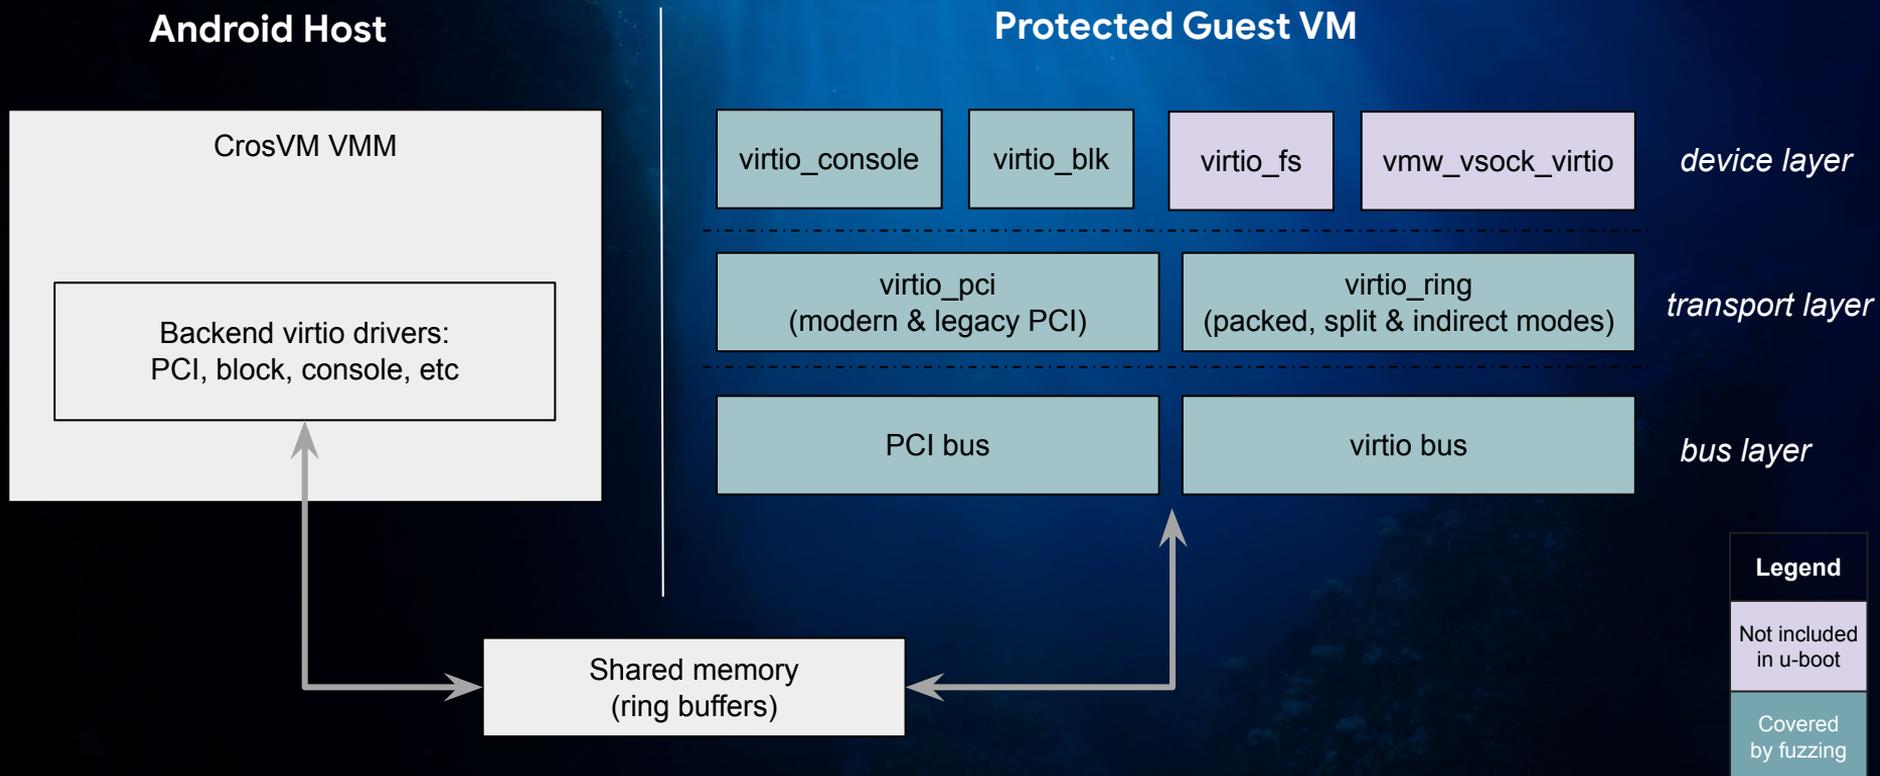
GEEKCON

# Virtio fuzzing for Android pKVM

# Android Protected KVM Attack Surface



# Protected VM virtio attack surface



# Virtio Front-end Fuzzers

## Kernel under test:

- android13-5.10

## virtio\_ring:

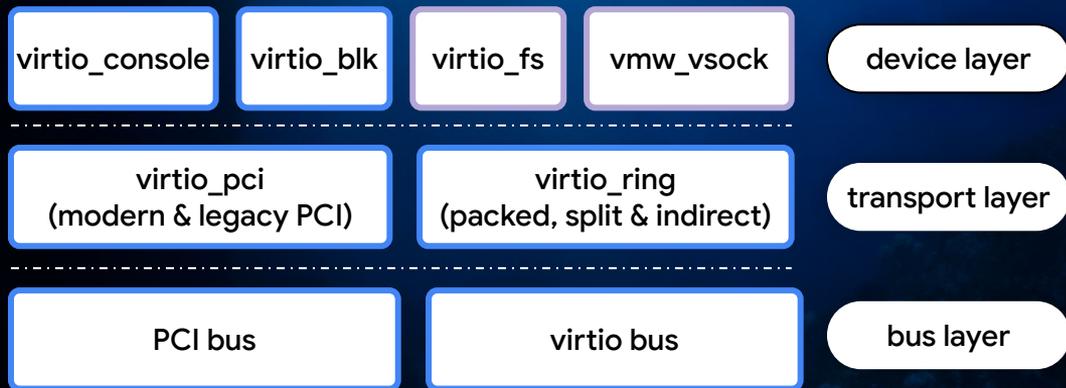
- fuzzes ring-buffer processing functionality
- handles both split & packed mode

## virtio\_pci:

- fuzzes PCI configuration space
- LKL arch-specific implementation of PCI bus
- mock-out PCI MMIO in the fuzzer harness

## virtio\_blk:

- mutates the virtio\_blk configuration block



# Virtio\_blk fuzzer finding

```
int block_read_full_page(struct page *page, get_block_t *get_block)
{
    struct buffer_head *bh, *head, *arr[MAX_BUF_PER_PAGE];

    ...
    do {
        if (buffer_uptodate(bh))
            continue;

        if (!buffer_mapped(bh)) {
            int err = 0;

            ...
        }

        arr[nr++] = bh;    <===== 00B write on stack

    } while (i++, iblock++, (bh = bh->b_this_page) != head);
    ...
}
```

# Virtio\_blk fuzzer finding

- With the block size **0xe5e5e5e5**:
  - ``inode->i_blkbits == 32``
  - ``1 << READ_ONCE(inode->i_blkbits)`` is undefined behavior in C
  - ``1 << READ_ONCE(inode->i_blkbits) == 1`` on x86 architecture

```
static struct buffer_head *create_page_buffers(struct page *page, ...)
{
    BUG_ON(!PageLocked(page));
    if (!page_has_buffers(page))
        create_empty_buffers(page, 1 << READ_ONCE(inode->i_blkbits), b_state);
    return page_buffers(page);
}
```

GEEKCON

DEMO

USB HID Driver issue discovery

GEEKCON

# Conclusion

# Key takeaways

- LKL can be an effective way to fuzz certain kernel scenarios
  - Scenarios requiring complex setup
  - Dependencies to some userspace libraries
  - Visual code coverage and debugging can be very useful
  - It's always good to have alternatives
- The fact that you don't know something might be your advantage
  - Discovering new approaches from naive ideas
  - Learning without predefined mindset

# Future work

- Upstream Binder fuzzer to LKL repo
- Adding MMU support
- Creating more fuzzers

# Acknowledgement

- Octavian Purdila and LKL maintainers
- <https://github.com/atrosinenko/kbdysch>
- Android pKVM team
- Android Binder team

GEEKCON



THANKS

Questions?

*androidoffsec-external@google.com*  
*<https://androidoffsec.withgoogle.com>*



# Resources

- [libFuzzer – a library for coverage-guided fuzz testing](#)
- [Achieving Kernel Code Execution through Malicious USB device](#)
- [Android Virtualization Framework and pKVM security](#)
- [Emulating USB device with Python](#)
- [Attacking Android Binder: Analysis and Exploitation of CVE-2023-20938](#)
- [Binder Internals](#)