

# Whole-System Dynamic Binary Analysis

Continued

# Repeatable Reverse Engineering for the Greater Good with PANDA

Brendan Dolan-Gavitt, Josh Hodosh, Patrick Hulin,  
Tim Leek, Ryan Whelan

## PANDA's Workflow




Fig. 1: Replay-based Reverse Engineering Workflow. PANDA's ability to record and replay whole system executions is the foundation of its use in reverse engineering. One captures a recording and then iteratively builds one or more plugins that perform dynamic analyses.

## Record&Replay Internals




Fig. 2: PANDA Non-determinism log

## Log sizes for various replays

Replay	Instructions	Log size	Instr/byte
frebsdboot.rr	9.3B	533MB	17
spotify.rr	12B	229MB	52
haikuurl.rr	8.6B	119MB	72
carberpl.rr	9.1B	43MB	212
win7iessl.rr	8.6B	9.4MB	915
Starcraft.rr	60M	1.8MB	33

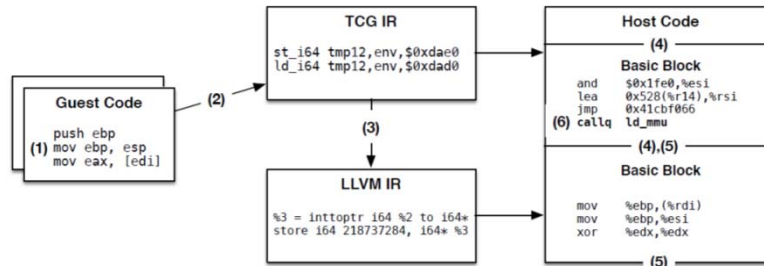
TABLE I: ND log sizes for various replays

## Record and Replay Runtime Overheads

Environment	Time in sec	Slowdown wrt Qemu 2.1.0
Qemu 2.1.0	35.6	1.0
PANDA	37.2	1.05
PANDA+record	66	1.85
PANDA+replay	127	3.57

TABLE II: PANDA, record, and replay slowdowns

## PANDA execution and instrumentation




## PANDA Plugins

- Tappan Zee (North) Bridge
- System Calls
- Shadow Callstack
- Taint Analysis
- Scissors



**DroidScope:**  
Seamlessly Reconstructing the OS and Dalvik Semantic Views for Dynamic Android Malware Analysis

Lok Yan  
Heng Yin  
August 10, 2012



**L.C. Smith** College of Engineering and Computer Science

# Android



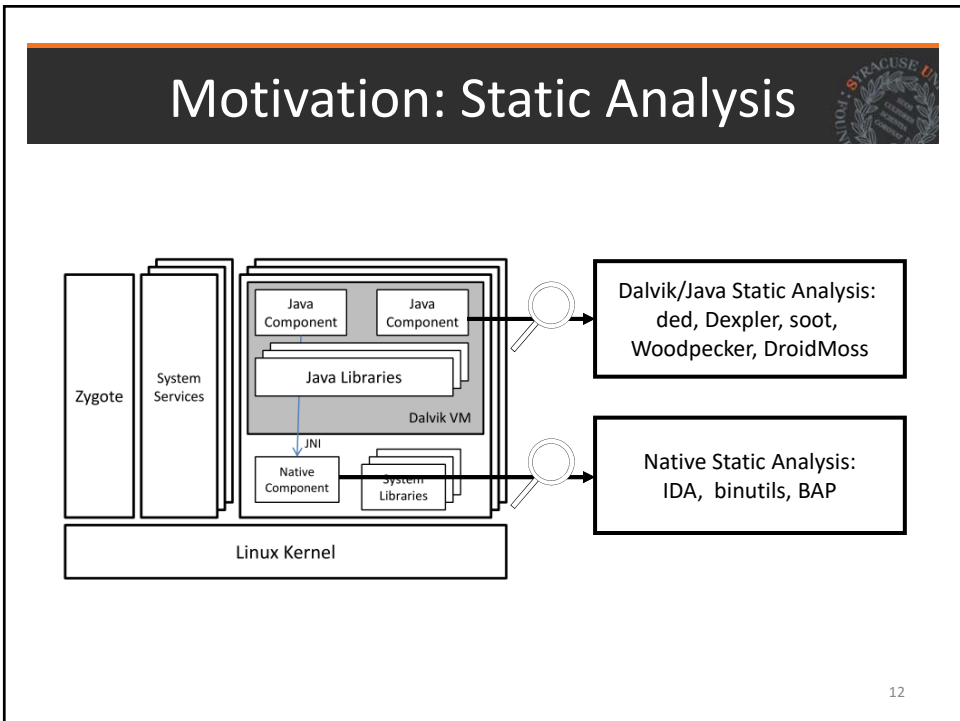
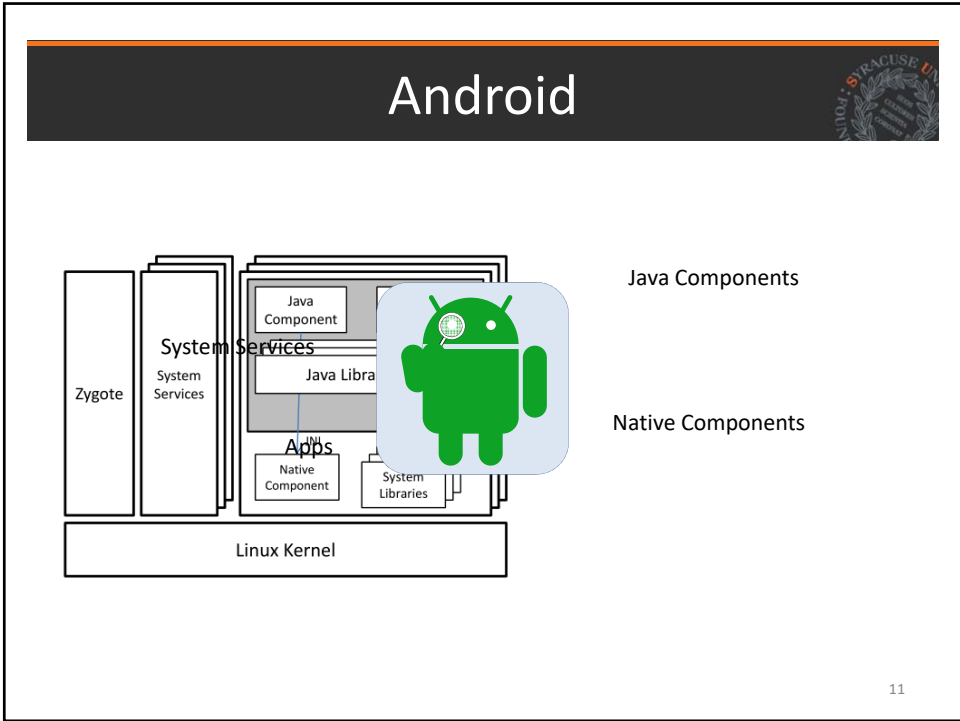
System Services

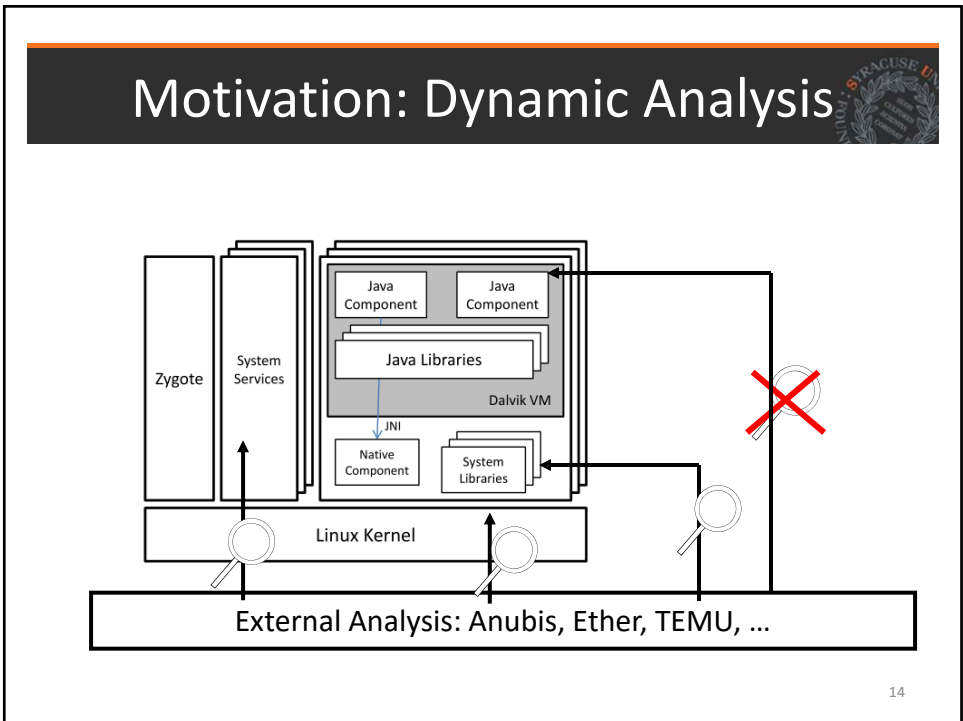
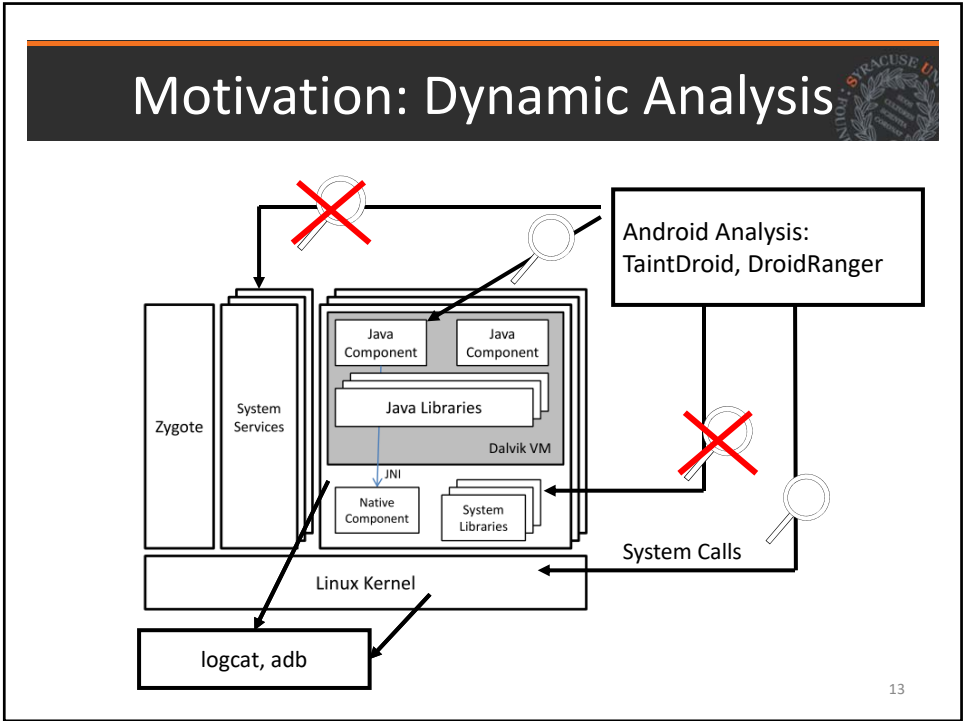
Apps

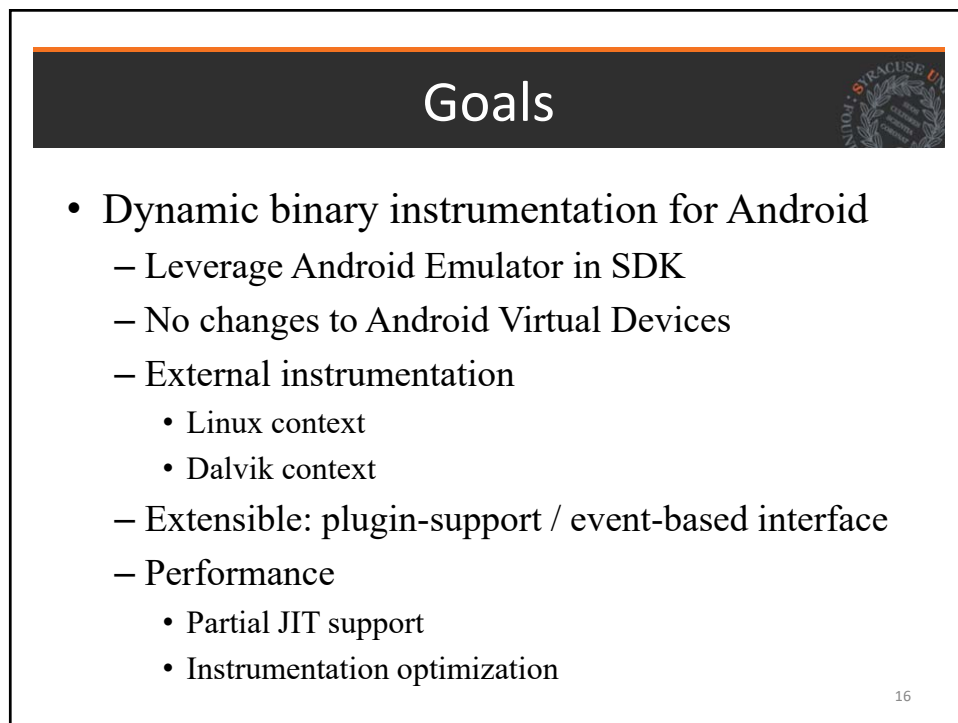
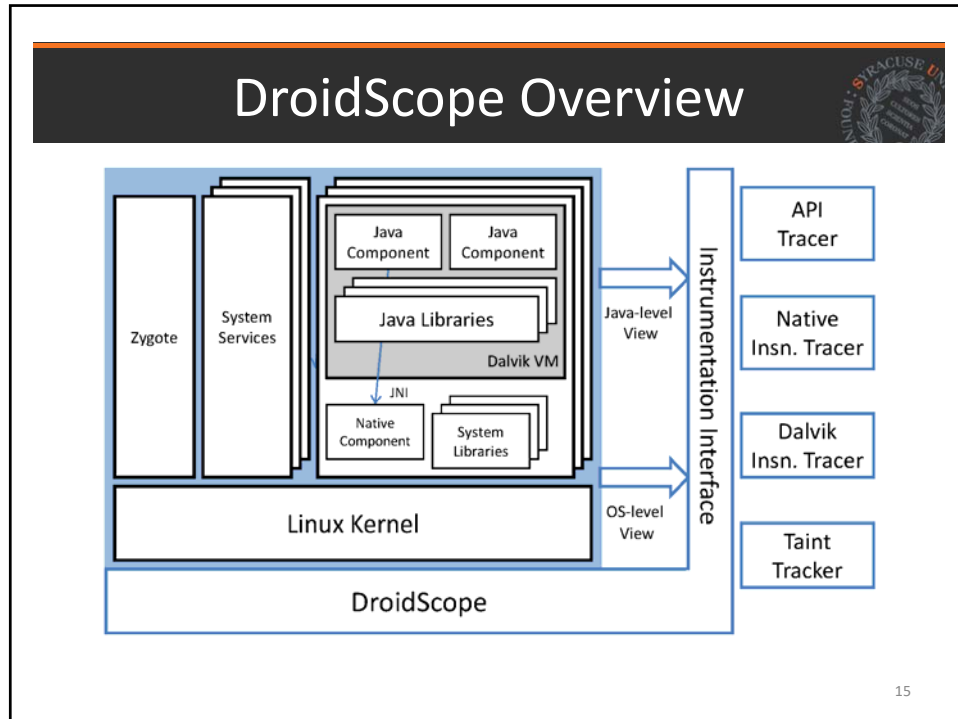
Java Components

Native Components

10









## Roadmap

- External instrumentation
  - Linux context
  - Dalvik context
- Extensible: plugin-support / event-based interface
- Evaluation
  - Performance
  - Usage

17

## Linux Context: Identify App(s)

- Shadow task list

```

lok@cypress: ~/temu_android/trunk/qemu/objs
File Edit View Terminal Help
PID  TGID Parent  UID   GID   COMM          PGD
323  323   32  10017 10017 com.cooliris.media 0x02ca0000
    TID, 6ThreadInfo
    323, 0xc2d60000
    324, 0xc40de000
    325, 0xc407e000
    326, 0xc401a000
    327, 0xc2d66000
    328, 0xc2ca6000
    329, 0xc4230000
    330, 0xc41d0000
    331, 0xc0e96000

PID  TGID Parent  UID   GID   COMM          PGD
333  333   32  10020 10020 com.android.calculator2 0x02158000
    TID, 6ThreadInfo
    333, 0xc5f12000
    334, 0xc05cc000
    335, 0xc256c000
    336, 0xc27f8000
    337, 0xc2570000
    338, 0xc21e0000
    339, 0xc327c000
    340, 0xc3292000
    342, 0xc2650000
(qemu)
  
```

## Java/Dalvik View

- Dalvik virtual machine
  - register machine (all on stack)
  - 256 opcodes
  - saved state, *glue*, pointed to by ARM R6, on stack in x86
- mterp
  - offset-addressing: *fetch opcode* then jump to  $(dvmAsmInstructionStart + opcode * 64)$
  - *dvmAsmSisterStart* for emulation overflow
- Which Dalvik opcode?
  1. Locate *dvmAsmInstructionStart* in shadow memory map
  2. Calculate  $opcode = (R15 - dvmAsmInstructionStart) / 64$ .

19

## Just In Time (JIT) Compiler

- Designed to boost performance
- Triggered by counter - mterp is always the default
- Trace based
  - Multiple basic blocks
  - Multiple exits or *chaining cells*
  - Complicates external introspection
  - Complicates instrumentation

20

## Disabling JIT

Update  
Program Counter(PC)

dvmGetCodeAddr(PC)  
 != NULL

21

## Roadmap

- ✓ External instrumentation
  - Linux context
  - Dalvik context
- Extensible: plugin-support / event-based interface
- Evaluation
  - Performance
  - Usage

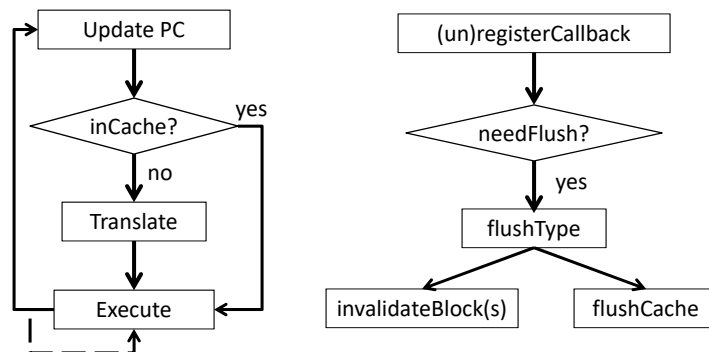
22

## Instrumentation Design

- Event based interface
  - Execution: e.g. native and Dalvik instructions
  - Status: updated shadow task list
- Query and Set, e.g. interpret and change cpu state
- Performance
  - Example: Native instructions vs. Dalvik instructions
  - Instrumentation Optimization

23

## Dynamic Instrumentation



24

# Instrumentation

	NativeAPI	LinuxAPI	DalvikAPI
Events	instruction begin/end	context switch	Dalvik instruction begin
	register read/write	system call	method begin
	memory read/write	task begin/end	
	block begin/end	task updated	
		memory map updated	
Query & Set	memory read/write	query symbol database	query symbol database
	memory r/w with pgd	get current context	interpret Java object
	register read/write	get task list	get/set DVM state
	taint set/check		taint set/check objects disable JIT

25

# Dalvik Instruction Tracer (Example)

```

1. void opcode_callback(uint32_t opcode) {
2.     printf("[%x] %s\n", GET_RPC, opcodeToStr(opcode));
3. }
4.
5. void module_callback(int pid) {
6.     if (bInitialized || (getIBase(pid) == 0))
7.         return;
8.
9.     getModAddr("dfk@classes.dex", &startAddr, &endAddr);
10.
11.     addDisableJITRange(pid, startAddr, endAddr);
12.     disableJITInit(getGetCodeAddrAddress(pid));
13.     addMterpOpcodesRange(pid, startAddr, endAddr);
14.     dalvikMterpInit(getIBase(pid));
15.     registerDalvikInsnBeginCb(&opcode_callback);
16.     bInitialized = 1;
17. }
18.
19. void _init() {
20.     setTargetByName("com.andhuhu.fengyinchuanshuo");
21.     registerTargetModulesUpdatedCb(&module_callback);
22. }

```

26

## Plugins



- API Tracer
  - System calls
    - *open, close, read, write*, includes parameters and return values
  - Native library calls
  - Java API calls
    - Java Strings converted to C Strings
- Native and Dalvik Instruction Tracers
- Taint Tracker
  - Taints ARM instructions
  - One bit per byte
  - Data movement & Arithmetic instructions including barrel shifter
  - Does not support control flow tainting

27

## Roadmap



- ✓ External instrumentation
  - Linux context
  - Dalvik context
- ✓ Extensible: plugin-support / event-based interface
- Evaluation
  - Performance
  - Usage

28

## Implementation



- Configuration
  - QEMU 0.10.50 – part of Gingerbread SDK
  - Gingerbread
    - “user-eng”
    - No changes to source
  - Linux 2.6.29, QEMU kernel branch

29


## Performance Evaluation



- Seven free benchmark Apps
  - AnTuTu Benchmark
  - (ABenchMark) by AnTuTu
  - CaffeineMark by Ravi Reddy
  - CF-Bench by Chainfire
  - Mobile processor benchmark (Multicore) by Andrei Karpushonak
  - Benchmark by Softweg
  - Linpack by GreeneComputing
- Six tests repeated five times each
  - Baseline
  - NO-JIT Baseline – uses a build with JIT disabled at runtime
  - Context Only
  - API Tracer
  - Dalvik Instruction Trace
  - Taint Tracker

30

## Select Performance Results



## Usage Evaluation

- Use DroidScope to analyze real world malware
  - API Tracer
  - Dalvik Instruction Tracer + dexdump
  - Taint Tracker – taint IMEI/IMSI @ *move\_result\_object* after *getIMEI/getIMSI*
- Analyze included exploits
  - Removed patches in Gingerbread
  - Intercept system calls
  - Native instruction tracer




## Droid Kung Fu

- Three encrypted payloads
  - ratc (Rage Against The Cage)
  - killall (ratc wrapper)
  - gjsvro (udev exploit)
- Three execution methods
  - piped commands to a shell (default execution path)
  - Runtime.exec() Java API (instrumented path)
  - JNI to native library terminal emulator (instrumented path)
  - Instrumented return values for *isVersion221* and *getPermission* methods

33

## Droid Kung Fu: TaintTracker



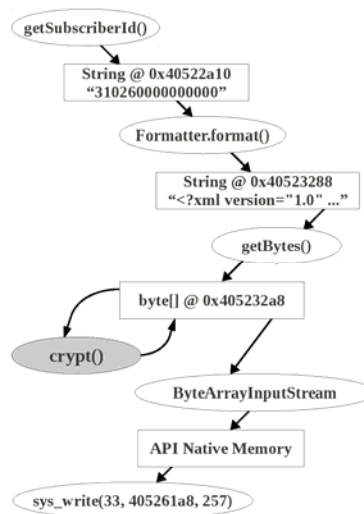
34

## DroidDream

- Same payloads as DroidKungFu
- Two processes
  - Normal *droiddream* process clears logcat
  - *droiddream:remote* is malicious
- xor-encrypts private information before leaking
- Instrumented *sys\_connect* and *sys\_write*

35

## Droid Dream: TaintTracker



36

## DroidDream: crypt trace



```
[43328f40] aget-byte v2(0x01), v4(0x405232a8), v0(186)
  Getting Tainted Memory: 40523372(2401372)
  Adding M@410acce(42c5cec) len = 4
[43328f44] sget-object v3(0x0000005e), KEYVALUE// field@0003
[43328f48] aget-byte v3(0x88), v3(0x4051e288), v1(58)
[43328f4c] xor-int/2addr v2(62), v3(41)
  Getting Tainted Memory: 410acce(42c5cec)
  Adding M@410acce(42c5cec) len = 4
[43328f4e] int-to-byte v2(0x17), v2(23)
  Getting Tainted Memory: 410acce(42c5cec)
  Adding M@410acce(42c5cec) len = 4
[43328f50] aput-byte v2(0x17), v4(0x405232a8), v0(186)
  Getting Tainted Memory: 410acce(42c5cec)
  Adding M@40523372(2401372) len = 1
```

37

## Summary



- DroidScope
  - Dynamic binary instrumentation for Android
  - Built on Android Emulator in SDK
  - External Introspection & Instrumentation support
  - Four plugins
    - API Tracer
    - Native Instruction Tracer
    - Dalvik Instruction Tracers
    - TaintTracker
  - Partial JIT support

38

## Related Works

- Static Analysis
  - ded, Dexpler, soot
  - Woodpecker, DroidMoss
- Dynamic Analysis
  - TaintDroid
  - DroidRanger
  - PIN, Valgrind, DynamoRIO
  - Anubis, TEMU, Ether, PinOS
- Introspection
  - Virtuoso
  - VMWatcher

39

## Challenges

- JIT
  - Full JIT support
  - Flushing JIT cache
- Emulation detection
  - Real Sensors: GPS, Microphone, etc.
  - Bouncer
- Timing assumptions, timeouts, events
- Closed source systems, e.g. iOS

40

# Questions?

Q0. Where can I get DroidScope?



L.C. Smith College of Engineering and Computer Science

## ratc



- Vulnerability
  - *setuid()* fails when `RLIMIT_NPROC` reached
  - *abdd* fails to verify *setuid()* success
- Three generation (stage) exploit
  - Locate *abdd* in */proc* and spawns child
  - Child *fork()* processes until `-11 (-EAGAIN)` is returned then spawns child – continues *fork()*
  - Grandchild *kill()* *abdd* and waits for process to re-spawn

## ratc: exploit diagnosis

```

///setgid returns from kernel back to adbd
0000813c: pop {r4, r7}
00008140: movs r0, r0
00008144: bxp1 lr : Read Oper[0]. R14, Val = 0xc3a5
// Return back to 0xc3a4 (caller) in Thumb mode

///adbd_main sets up for setuid
0000c3a4: movs r0, #250
0000c3a6: lsls r0, r0, #3 : Write Oper[0]. R0, Val = 0x7d0
// 250 * 8 = 0x7d0 = 2000 = AID_SHELL
...
///Start of setuid section
// 213 is syscall number for sys_setuid
000082e0: push {r4, r7} : Write Oper[0]. M8be910bb8, Val = 0x7d0
// push AID_SHELL onto the stack
000082e4: mov r7, #213
000082e8: svc 0x00000000
// Make sys call

// == TRANSITION TO KERNEL SPACE ==
///sys_setuid then calls set_user in kernel mode
///inside sys_setuid
// Has rlimit been reached?
c0048944: cmp r2, r3 : Read Oper[0]. R3, Val = 300 Read Oper[1]. R2, Val = 300
// RLIMIT(300) is reached and !init_user so return -11
c0048960: mvn r0, #10 : Write Oper[0]. R0, Val = 0xffffffff5
// the return value is now -11 or -EAGAIN
c0048964: ldmb sp, {r4, r5, r6, fp, sp, pc}
///Return back to sys_setuid which returns back to userspace

// == RETURN TO USERSPACE ==
///setuid continues
000082ac: pop {r4, r7}
000082d0: movs r0, r0 : Read Oper[0]. R0, Val = 0xffffffff5
// -11 is still here
///Return back to adb_main at 0xc3ac (the return address) above
// Immediately starts other work, does not check return code
0000c3ac: ldr r7, [pc, #356] : Read Oper[0]. M8000c514, Val = 0x19980330
Write Oper[0]. R7, Val = 0x19980330
// 0x19980330 is _LINUX_CAPABILITY_VERSION

```

43

## Symbol Information

- Native library symbols - Static
  - From *objdump* of libraries
- Java symbols - Dynamic
  - Dalvik data structures -> address of string
  - Given address, load from
    - Memory
    - File mapped into memory
  - *dexdump* as backup

44

