

# eBPF - Android Reverse Engineering Superpowers

Terry Chia

## \$ whoami

- Security Consultant @ Centurion Information Security
- Blog: <https://www.ayrx.me>
- Github: <https://github.com/Ayrx>

# Agenda

- What is eBPF?
- Why eBPF?
- eBPF on Android
- Writing eBPF programs

# What is eBPF?

# eBPF

- extended Berkeley Packet Filter
- Framework for tracing a Linux system
- Linux version of Solaris' DTrace
- Requires a relatively new Linux kernel (> 4.1). Newer kernels might have more features.

# eBPF

- Write custom code that triggers whenever *something* happens in the system.
- Write eBPF program & compile to bytecode.
- eBPF bytecode is loaded through the *bpf(2)* syscall.
- eBPF code is executed with an in-kernel virtual machine.

# eBPF

- Event Sources:
  - kprobes / kretprobes
  - uprobes / uretprobes
  - Tracepoints
  - User Statically-Defined Tracing Probes (USDT)

## k(ret)probes, u(ret)probes

- kprobes / kretprobes are used to attach to kernel functions.
- uprobes / uretprobes are used to attach to userspace functions.



# Tracepoints

- Tracepoints are used to attach to *events* within the kernel.
- A large number of events are exposed by the kernel. See [/sys/kernel/debug/tracing/events/](#) for a full list.

# Tracepoints

```
root@dlvisi0n:~# ls /sys/kernel/debug/tracing/events/
alarmtimer  filemap      irq_matrix   module       regmap       thermal_power_allocator
asoc        fs           irq_vectors  mpx          regulator    timer
block       fs_dax       iwlwifi      msr          rpm          tlb
bridge      ftrace       iwlwifi_data napi          rseq         ucspi
btrfs       gpio         iwlwifi_io   net          rtc          udp
cfg80211    gvt          iwlwifi_msg  nfsd         sched        v4l2
cgroup      hda          iwlwifi_ucode nmi          scsi         vb2
clk         hda_controller jbd2         nvme         oom          vmscan
cma         hda_intel   kvm          page_isolation skb          vsyscall
compaction  header_event kmmmu        pagemap      sock         workqueue
cpuhp       header_page libata        percpu       spi          writeback
dma_fence   huge_memory mac80211      power        sunrpc       x86_fpu
drm         hyperv      mac80211_msg printk        swiotlb      xdp
enable      i2c         mce           random       sync_trace   xen
exceptions  i915        mdio          ras          syscalls     xhci-hcd
ext4        initcall    mei           raw_syscalls task          thermal
fib         intel-sst   migrate       rcu
fib6        iommu       mmc
filelock    irq
```

root@dlvisi0n:~# █

# User Statically-Defined Tracing Probes (USDT)

- USDTs are Tracepoints for userspace.

```
1  #include <sys/sdt.h>
2  #include <sys/time.h>
3  #include <unistd.h>
4
5  int main(int argc, char **argv)
6  {
7      struct timeval tv;
8
9      while(1) {
10         gettimeofday(&tv, NULL);
11         DTRACE_PROBE1(test-app, test-probe, tv.tv_sec);
12         sleep(1);
13     }
14     return 0;
15 }
```

# eBPF

- eBPF also provides a way for userspace to communicate with a eBPF program.
- BPF\_PERF\_OUTPUT
  - Fast ring buffer
  - Can create multiple ring buffers per eBPF program

# eBPF

- Don't write eBPF bytecode by hand.
- Use the *bcc* compiler!
  - <https://github.com/iovisor/bcc>

# eBPF

- Running arbitrary code in the kernel is risky.
- eBPF has a validator that tries to ensure that eBPF programs are “safe”
  - eBPF program must terminate
  - Validates stack / register state
  - Validates no out-of-bounds reads

# Current eBPF Usage

- eBPF is mainly used for instrumenting *production* Linux systems.
- Especially popular in container / kubernetes environment.
- Firewalls -  
<https://cilium.io/blog/2018/11/20/fb-bpf-firewall/>

# Why eBPF?



# eBPF - Android Reverse Engineering Superpowers

# Reverse Engineering

- Reverse Engineering is about understanding an application.
- Three main categories of techniques:
  - Static Analysis - IDA Pro / Ghidra
  - Debugging - GDB / WinDBG / Intel PIN / Frida
  - Behavioural Analysis - strace / ltrace / Procmon

# Anti-Reversing

- Anti-Reversing tricks for each technique:
  - Static Analysis - Obfuscation
  - Debugging - Anti-debugging, Root / Jailbreak Detection
  - Behavioural Analysis - Anti-debugging, Root / Jailbreak Detection

# Anti-Reversing

- OWASP MSTG describes some common anti-reversing techniques:
  - <https://mobile-security.gitbook.io/mobile-security-testing-guide/android-testing-guide/0x05j-testing-resiliency-against-reverse-engineering>
  - <https://mobile-security.gitbook.io/mobile-security-testing-guide/ios-testing-guide/0x06j-testing-resiliency-against-reverse-engineering>

# Android Anti-Reversing

- Android applications commonly utilize a combination of the following tricks:
  - Root Detection
  - Anti-Debugging
  - Obfuscation
  - File Integrity Checks

# Android Anti-Reversing

- Tricks that prevent *ptrace* from being used are particularly annoying.
- Many common tools rely on being able to *ptrace* the target process.
  - strace
  - Gdb
  - Frida

## Anti-ptrace Example

- Only one tracer can be attached to a process.
- *ptrace*-ing a process that is already being debugged by another process will fail.
- A common technique is for an application to fork a child process that attaches to the parent process.

# Anti-ptrace Example

```
1 void fork_and_attach()
2 {
3     int pid = fork();
4
5     if (pid == 0)
6     {
7         int ppid = getppid();
8
9         if (ptrace(PTRACE_ATTACH, ppid, NULL, NULL) == 0)
10        {
11            waitpid(ppid, NULL, 0);
12
13            /* Continue the parent process */
14            ptrace(PTRACE_CONT, NULL, NULL);
15        }
16    }
17 }
```

- Fork a child process.



# Anti-pttrace Example

```
1 void fork_and_attach()
2 {
3     int pid = fork();
4
5     if (pid == 0)
6     {
7         int ppid = getppid();
8
9         if (ptrace(PTRACE_ATTACH, ppid, NULL, NULL) == 0)
10        {
11            waitpid(ppid, NULL, 0);
12
13            /* Continue the parent process */
14            ptrace(PTRACE_CONT, NULL, NULL);
15        }
16    }
17 }
```

- Fork a child process.
- Child process calls *ptrace* on the parent process.
- As long as the child process is alive, no other process can *ptrace* the parent.

# Anti-ptrace Example

- In the trivial example, killing the child process will allow the parent process to be debugged.
- In practice, tricks to ensure that the child process remains alive can be used.
  - Forking multiple processes tracing each other
  - Monitoring running processes

# Android Anti-Reversing

- Applications with multiple anti-RE tricks implemented can be difficult to analyze.
- Low level (kernel!) capabilities are really helpful to debug such applications.
- You essentially want capabilities at a level of the system that the application cannot subvert.

# Android Anti-Reversing

- Previously, this meant writing a kernel module.
  - Error-prone: Bad code means crashing your device
  - Tedious development process: Write -> Compile -> Transfer to Device -> Hope it works
- What other options do we have if we want to run custom code in the kernel? 🤔

# eBPF on Android

# Requirements & Setup

- eBPF works on Android because Android uses a relatively standard Linux kernel.
- Took me a while to figure out how to get everything working.
- Most of the documentation for the eBPF toolchain assumes standard Linux instead of Android.

# bcc toolchain

- *bcc* is the standard eBPF compiler toolchain.
  - <https://github.com/iovisor/bcc>
- It is a LLVM-based compiler toolchain that compiles C code to eBPF bytecode.
- Requires kernel headers to be present.
- Requires Python.

# adeb

- *adeb* makes it easy to setup the bcc toolchain.
  - <https://github.com/joelagnel/adeb>
- Essentially a Debian-based environment running on the Android device via chroot magic.
- Comes with *bcc* and other useful tools.



# Building the Android Kernel

- *adep* still requires a kernel with the required configs turned on.
- Kernel version 4.9 and above
  - CONFIG\_KPROBES=y
  - CONFIG\_KPROBE\_EVENT=y
  - CONFIG\_BPF\_SYSCALL=y
  - CONFIG\_IKHEADERS=m
  - CONFIG\_UPROBES=y
  - CONFIG\_UPROBE\_EVENT=y

# Building the Android Kernel

- No current Android ships with the necessary configs.
  - Have not looked at Android Q yet.
- This means building your own kernel.
- Process differs depending on where you are running Android (Emulator vs device, etc)

# adeb

- `adeb prepare --full --build --kernelsrc path/to/kernel`
  - `--arch <amd64/arm64/etc>`
  - `--builddir <output dir>`
- `adeb prepare --archive <output dir>/androdeb-fs.tgz --kernelsrc path/to/kernel`

# adeb

- adeb shell
- *bcc* by default comes with some really useful utilities.
  - filetop
  - opensnoop (and other \*snoop commands)
  - etc



# Writing eBPF programs

## bcc toolchain

- With *bcc*, you write eBPF programs in Python.
  - `import bcc`
  - Build a string that contains the eBPF program.
  - Pass string to `bcc` which invokes LLVM behind the scenes and loads the compiled program into the kernel.

# Hello World

- Write an eBPF program that prints out all the files opened on the system.
  - Attach to an appropriate kernel function
  - Send the pathname being opened to userspace
  - Print the output



# Hello World

```
1  #!/usr/bin/python
2
3  from bcc import BPF
4
5  program = """
6  #include <asm/ptrace.h>
7  #include <uapi/linux/limits.h>
8
9  struct open_data_t {
10     char fname[NAME_MAX];
11 };
12
13 BPF_PERF_OUTPUT(open_event);
14
15 int kprobe__sys_openat(struct pt_regs *ctx,
16     int dirfd, char __user* pathname, int flags, mode_t mode) {
17
18     struct open_data_t data = {};
19     bpf_probe_read(&data.fname, sizeof(data.fname), (void *)pathname);
20     open_event.perf_submit(ctx, &data, sizeof(data));
21
22     return 0;
23 }
24 """
```

- *open\_data\_t* is a struct that stores data we want to send to userspace.
- BPF\_PERF\_OUTPUT opens up a ring buffer called *open\_event*.

# Hello World

```
1  #!/usr/bin/python
2
3  from bcc import BPF
4
5  program = """
6  #include <asm/ptrace.h>
7  #include <uapi/linux/limits.h>
8
9  struct open_data_t {
10 |     char fname[NAME_MAX];
11 | };
12
13  BPF_PERF_OUTPUT(open_event);
14
15  int kprobe_sys_openat(struct pt_regs *ctx,
16 |     int dirfd, char __user* pathname, int flags, mode_t mode) {
17
18 |     struct open_data_t data = {};
19 |     bpf_probe_read(&data.fname, sizeof(data.fname), (void *)pathname);
20 |     open_event.perf_submit(ctx, &data, sizeof(data));
21
22 |     return 0;
23 | }
24  """
```

- `kprobe_` syntax tells `bcc` that the function is a kprobe.
- The function arguments to the syscall can be omitted if your eBPF program does not use them.

# Hello World

```
1  #!/usr/bin/python
2
3  from bcc import BPF
4
5  program = """
6  #include <asm/ptrace.h>
7  #include <uapi/linux/limits.h>
8
9  struct open_data_t {
10 |     char fname[NAME_MAX];
11 | };
12
13  BPF_PERF_OUTPUT(open_event);
14
15  int kprobe_sys_openat(struct pt_regs *ctx,
16 |     int dirfd, char __user* pathname, int flags, mode_t mode) {
17
18 |     struct open_data_t data = {};
19 |     bpf_probe_read(&data.fname, sizeof(data.fname), (void *)pathname);
20 |     open_event.perf_submit(ctx, &data, sizeof(data));
21
22 |     return 0;
23 | }
24  """
```

- Initialize an instance of *open\_data\_t* to store the file name.
- Use the special *bpf\_probe\_read* function to copy the data into the *fname* array.
- *open\_event.perf\_submit* sends the initialized *open\_data\_t* instance to userspace.

# Hello World

```
25
26 def print_open_event(cpu, data, size):
27     event = b["open_event"].event(data)
28     print event.fname
29
30 b = BPF(text=program)
31 b["open_event"].open_perf_buffer(print_open_event)
32
33 while True:
34     try:
35         b.perf_buffer_poll()
36     except KeyboardInterrupt:
37         exit()
38
```

- *print\_open\_event* is a callback function that can be made to trigger when a perf event is received.

# Hello World

```
25
26 def print_open_event(cpu, data, size):
27     event = b["open_event"].event(data)
28     print event.fname
29
30 b = BPF(text=program)
31 b["open_event"].open_perf_buffer(print_open_event)
32
33 while True:
34     try:
35         b.perf_buffer_poll()
36     except KeyboardInterrupt:
37         exit()
38
```

- *print\_open\_event* is a callback function that can be made to trigger when a perf event is received.
- Initialize an eBPF program and opens up the perf buffer called *open\_event*.

# Hello World

```
25
26 def print_open_event(cpu, data, size):
27     event = b["open_event"].event(data)
28     print event.fname
29
30 b = BPF(text=program)
31 b["open_event"].open_perf_buffer(print_open_event)
32
33 while True:
34     try:
35         b.perf_buffer_poll()
36     except KeyboardInterrupt:
37         exit()
38
```

- *print\_open\_event* is a callback function that can be made to trigger when a perf event is received.
- Initialize an eBPF program and opens up the perf buffer called *open\_event*.
- Polls all opened buffers in an infinite loop.



# Code Generation

- Codegen is a common pattern that you will see in eBPF programs.
- Useful if there is a part of the code you want to change every time you run the program.
  - PID filtering is one example.
- Codegen is also useful to get around eBPF limitations.



# Code Generation

```
1  #!/usr/bin/python
2
3  from bcc import BPF
4
5  program = """
6  #include <asm/ptrace.h>
7  #include <uapi/linux/limits.h>
8
9  struct open_data_t {
10     char fname[NAME_MAX];
11 };
12
13 BPF_PERF_OUTPUT(open_event);
14
15 int kprobe__sys_openat(struct pt_regs *ctx,
16     int dirfd, char __user* pathname, int flags, mode_t mode) {
17
18     PID_FILTER
19
20     struct open_data_t data = {};
21     bpf_probe_read(&data.fname, sizeof(data.fname), (void *)pathname);
22     open_event.perf_submit(ctx, &data, sizeof(data));
23
24     return 0;
25 }
26 """
```

# Code Generation

```
33 def insert_pid_filter(bpf_text, pid):
34     bpf_text = "#define FILTER_PID {}\n".format(pid) + bpf_text
35     pid_filter = ""
36     u64 pid_tgid = bpf_get_current_pid_tgid();
37     if (pid_tgid >> 32 != FILTER_PID) {
38         return 0;
39     }
40     ""
41     bpf_text = bpf_text.replace("PID_FILTER", pid_filter)
42
43     return bpf_text
44
45 program = insert_pid_filter(program, sys.argv[1])
```

- *insert\_pid\_filter* replaces the `PID_FILTER` placeholder in the eBPF program string with C code.

# Code Generation

```
33 def insert_pid_filter(bpf_text, pid):
34     bpf_text = "#define FILTER_PID {}\n".format(pid) + bpf_text
35     pid_filter = ""
36     u64 pid_tgid = bpf_get_current_pid_tgid();
37     if (pid_tgid >> 32 != FILTER_PID) {
38         return 0;
39     }
40     ""
41     bpf_text = bpf_text.replace("PID_FILTER", pid_filter)
42
43     return bpf_text
44
45 program = insert_pid_filter(program, sys.argv[1])
```

- *insert\_pid\_filter* replaces the `PID_FILTER` placeholder in the eBPF program string with C code.
- The value of the `FILTER_PID` macro depends on the value of `sys.argv[1]`.



## strace.py

- Syscall tracing utility implemented with eBPF
- Trace mode vs Aggregate mode
- Filter by PID / Process Name
- Filter only syscalls you are interested in
- Disclaimer: Pretty ugly code



# Modifying the system with eBPF

- eBPF can *write* to userspace memory with the *bpf\_probe\_write\_user* function.
  - `int bpf_probe_write_user(void *dst, const void *src, u32 len)`
- This only works for userspace memory that already has write permissions in place.
  - So no writing to the `.text` segment with eBPF
  - You can however write to the stack, heap, etc

# Modifying the system with eBPF

- As an example, we can use this capability to bypass simple root detection techniques.
- OWASP MSTG - UnCrackable-Level3.apk
- Looks for the presence of certain files on the system.



# Modifying the system with eBPF

<code>/sbin/su</code>	<code>/system/app/Superuser.apk</code>
<code>/system/sbin/su</code>	<code>/system/xbin/daemonsu</code>
<code>/apex/com.android.runtime/bin/su</code>	<code>/system/etc/init.d/99SuperSUDaemon</code>
<code>/system/bin/su</code>	<code>/system/bin/.ext/.su</code>
<code>/system/xbin/su</code>	<code>/system/etc/.has_su_daemon</code>
<code>/odm/bin/su</code>	<code>/system/etc/.installed_su_daemon</code>
<code>/vendor/bin/su</code>	<code>/dev/com.koushikdutta.superuser.daemon</code>
<code>/vendor/xbin/su</code>	

# Modifying the system with eBPF

- *faccessat* syscall is used to test for the presence of the files.
  - `int faccessat(int dirfd, const char *pathname, int mode, int flags);`
- We can use *bpf\_probe\_write\_user* to modify the value of the *pathname* parameter.
- Redirect all *faccessat* calls on those files to a non-existent file.



# Closing Notes

# Limitations

- The verifier can get in the way of writing complex or substantial programs.
- Requires familiarity with Linux Kernel APIs.
- No stable kernel API. An eBPF program working on one version might break after an upgrade.
  - <https://github.com/torvalds/linux/blob/master/Documentation/process/stable-api-nonsense.rst>

# Limitations

- An eBPF program cannot write to kernel memory.
- TOCTOU issues are a huge problem when hooking syscalls.
  - Exploiting races in system call wrappers - <https://lwn.net/Articles/245630/>

## Awesome-ness

- eBPF offers a lot of power while being relatively simple to write.
  - Good for ad-hoc tracing or scripting
- Want to see: a set of eBPF programs like the ones in *bcc* but focused on security / RE.

# Useful Resources

- *bcc* Reference Guide -  
[https://github.com/iovisor/bcc/blob/master/docs/reference\\_guide.md](https://github.com/iovisor/bcc/blob/master/docs/reference_guide.md)
- *man* and other Linux kernel programming references
- BPF Performance Tools
  - <http://www.brendangregg.com/bpf-performance-tools-book.html>



Code:

<https://github.com/CenturionInfoSec/ebpf-examples>

Slides:

<https://bit.ly/2kUnlrg>







