

# Introduction to Embedded Software

## Practice #9

Kernel debugging mechanisms

**Dongkun Shin**

Embedded Software Laboratory

Sungkyunkwan University

<http://nyx.skku.ac.kr/>

# Objective

- Kernel debugging mechanisms
  - Debugfs
  - strace
  - ftrace
- Day practice
  - Day Practice: submit to i-campus (- 06/04 23:59)

# debugfs

- Debugfs exists as a simple way for kernel developers to make information available to user space
- Debugfs has no rules at all.
  - /proc: only meant for information about a process
  - Sysfs: strict one-value-per-file rules

```
root@raspberrypi:/sys/kernel/debug# ls
bcm2708_fb      brcmfmac      f2fs           ieee80211     mmci           regmap         tracing
bcm2835_thermal cleancache    fault_around_bytes irq            pinctrl       regulator      usb
bdi            clk           frontswap     kprobes       pm_genpd      sched_debug    vchiq
block         dma_buf       gpio          memblock      pm_qos        sched_features vc-mem
bluetooth     extfrag       hid           mmc0          pwm           sleep_time     vc-smem
```

```
root@raspberrypi:/sys/kernel/debug# cat usb/devices
```

```
T:  Bus=01 Lev=00 Prnt=00 Port=00 Cnt=00 Dev#=  1 Spd=480  MxCh=  1
B:  Alloc=  0/800 us ( 0%), #Int=  0, #Iso=  0
D:  Ver= 2.00 Cls=09(hub  ) Sub=00 Prot=01 MxPS=64 #Cfgs=  1
P:  Vendor=1d6b ProdID=0002 Rev= 4.14
S:  Manufacturer=Linux 4.14.98-v7+ dwc_otg_hcd
```

# Debugfs program

- **ext2/ext3/ext4 file system debugger**
  - Default filesystem in linux
- **Features**
  - Print file system information,
  - physical block,
  - inode information, path
  - <https://linux.die.net/man/8/debugfs>

```
debugfs: ncheck 20253
Inode    Pathname
20253    /home/pi/evtest.c
```

```
root@raspberrypi:/home/pi# debugfs
debugfs 1.43.4 (31-Jan-2017)
debugfs: open /dev/mmcblk0p2
debugfs: stats
Filesystem volume name:   rootfs
Last mounted on:         /
Filesystem UUID:         29075e46-f0d4-44e2-a9e7-
Filesystem magic number: 0xEF53
Filesystem revision #:   1 (dynamic)
Filesystem features:     has_journal ext_attr res
Filesystem flags:        signed_directory_hash
Default mount options:   user_xattr acl
Filesystem state:        clean
Errors behavior:         Continue
Filesystem OS type:      Linux
Inode count:              942480
Block count:              3877248
Reserved block count:    158905
Free blocks:              3495203
Free inodes:              895835
First block:              0
Block size:               4096
Fragment size:           4096
Reserved GDT blocks:     108
Blocks per group:        32768
Fragments per group:     32768
Inodes per group:        7920
Inode blocks per group:  495
Flex block group size:   16
```



# ftrace

- Ftrace brings to Linux is the ability to see what is happening inside the kernel.

- **Start ftrace**

```
$ sudo su
```

```
# get permission
```

```
# cd /sys/kernel/debug/tracing
```

```
# echo 1 > tracing_on
```

```
# ftrace enable
```

- **Set Event**

- Method 1: echo 1 > event/.../enable

- ✓ Example) 

```
# echo 1 > event/sched/enable
```

```
# scheduler profiling
```

- ✓ Get setting event list: # cat set\_event

```
root@raspberrypi:/sys/kernel/debug/tracing# cat set_event
root@raspberrypi:/sys/kernel/debug/tracing# echo 1 > events/sched/enable
root@raspberrypi:/sys/kernel/debug/tracing# cat set_event
sched:sched_wake_idle_without_ipi
sched:sched_swap_numa
sched:sched_stick_numa
sched:sched_move_numa
```

# ftrace

- **Set Event (cont`)**

- Method 2: echo "events" > set\_event
  - ✓ Get available event list: # cat available\_events
  - ✓ Example) # echo "sched" > set\_event # scheduler profiling
  - ✓ Example) # echo "sched\_switch" > set\_event # context switch profiling

- **Disable Event**

- Method 1: echo 0 > event/.../enable
  - ✓ Example) # echo 0 > event/sched/enable # Disable scheduler profiling
- Method 2: echo "" > set\_event # Disable all profiling

```
root@raspberrypi:/sys/kernel/debug/tracing# echo "" > set_event
root@raspberrypi:/sys/kernel/debug/tracing# cat set_event
root@raspberrypi:/sys/kernel/debug/tracing# echo "sched_switch" > set_event
root@raspberrypi:/sys/kernel/debug/tracing# cat set_event
sched:sched_switch
```

# ftrace

- **print ftrace**
  - Method 1: `# cat trace`
  - Method 2: `# cat trace_pipe` (live tracing)

```
root@raspberrypi:/sys/kernel/debug/tracing# cat trace | head -n 20
# tracer: nop
#
#          -----> irqs-off
#          /-----> need-resched
#          | /-----> hardirq/softirq
#          || /-----> preempt-depth
#          ||| /-----> delay
#
#          TASK-PID   CPU#   |      |      |      |      |      |
#          |  |      |   |      |      |      |      |      |
# migration/0-10    [000] d... 1220106.348201: sched_switch: prev_comm=migration/0 prev_pid=1
# migration/2-18    [002] d... 1220106.348201: sched_switch: prev_comm=migration/2 prev_pid=1
# migration/1-13    [001] d... 1220106.348201: sched_switch: prev_comm=migration/1 prev_pid=1
#          bash-14134 [003] d... 1220106.349626: sched_switch: prev_comm=bash prev_pid=14134 pr
# kworker/u8:0-14168 [003] d... 1220106.349653: sched_switch: prev_comm=kworker/u8:0 prev_pid=
#          <idle>-0    [000] d... 1220106.349670: sched_switch: prev_comm=swapper/0 prev_pid=0 p
```

- **Clearing the trace**
  - `# echo > trace`



# Day practice

- Insert “ftrace printk message” to module init function
  - `trace_printk(“[Your Name]\n”);`
  - Rebuild your module and copy module to raspberry device
- Set **only “module” event**
- Test Command

```
root@raspberrypi:/home/pi# sudo modprobe input-polldev
root@raspberrypi:/home/pi# insmod nunchuk.ko
root@raspberrypi:/home/pi# rmmod nunchuk
```

- Submit a “Screenshot” of “cat trace”

Day Practice: submit to i-campus (- 06/04 23:59)

# Day practice

- Submission Example

```
root@raspberrypi:/sys/kernel/debug/tracing# cat trace
# tracer: nop
#
#          -----=> irqs-off
#         /-----=> need-resched
#        | /-----=> hardirq/softirq
#       || /-----=> preempt-depth
#      ||| /-----=> delay
#
# TASK-PID  CPU#  ||||  TIMESTAMP  FUNCTION
#   | |   |   |   |   |   |
sudo-14499 [002] .... 1221201.384648: module_request: net-pf-16-proto-9 wait=1 call_site=netlink_create
sudo-14499 [002] .... 1221201.390287: module_request: net-pf-16-proto-9 wait=1 call_site=netlink_create
sudo-14499 [002] .... 1221201.395674: module_request: net-pf-16-proto-9 wait=1 call_site=netlink_create
sudo-14499 [002] .... 1221201.408309: module_request: net-pf-16-proto-9 wait=1 call_site=netlink_create
sudo-14499 [002] .... 1221201.413116: module_request: net-pf-16-proto-9 wait=1 call_site=netlink_create
sudo-14499 [002] .... 1221201.417611: module_request: net-pf-16-proto-9 wait=1 call_site=netlink_create
insmod-14507 [001] .... 1221203.704352: module_get: input_polldev call_site=ref_module refcnt=2
insmod-14507 [001] .... 1221203.704864: module_load: nunchuk 0
insmod-14507 [001] .... 1221203.704869: 0x7f0db000: ESLAB
insmod-14507 [001] .... 1221203.711519: module_put: nunchuk call_site=do_init_module refcnt=1
systemd-udevd-14508 [002] .... 1221203.732326: module_get: evdev call_site=chrdev_open refcnt=2
th-cmd-14511 [001] .... 1221203.745539: module_get: joydev call_site=chrdev_open refcnt=2
th-cmd-14511 [001] .... 1221203.840545: module_put: joydev call_site=__fput refcnt=1
systemd-udevd-14508 [002] .... 1221203.880573: module_put: evdev call_site=__fput refcnt=1
th-cmd-14512 [001] .... 1221203.894649: module_get: evdev call_site=chrdev_open refcnt=2
th-cmd-14512 [001] .... 1221203.970539: module_put: evdev call_site=__fput refcnt=1
rmmod-14513 [000] .... 1221205.691037: module_free: nunchuk
rmmod-14513 [000] .... 1221205.691208: module_put: input_polldev call_site=module_unload_free refcnt=1
```