

Assignment 3: Introduction to Linux Kernel Modules (Tutorial) (v0.10)

Note: Use of Ubuntu 24.04 is required for this assignment.

The purpose of this assignment is to introduce Linux Kernel programming in the context of writing a Linux Kernel module. Kernel code executes in kernel (privileged mode) where program instructions can command the CPU to perform any operation allowed by the CPU's architecture. This allows any instruction to be executed, any I/O operation to be initiated, and any area of memory accessed directly.

The host operating system is the base operating system of the computer which is started on boot time. The guest operating system is the operating system of any virtual machine that runs on the host. It is best to complete this tutorial on a virtual machine. In case of problems, it is easier to restart a virtual machine, than to restart the host system.

Assignment Submission

This assignment is accompanied by an online Canvas quiz which captures the answers for the questions in this assignment. After completing the activities described below, to complete the assignment and receive a grade, log into Canvas, and complete the Assignment #3 Quiz.

For this assignment, several packages for building kernel modules must be installed. The assumption is that the following packages have already been installed. To check if a package is installed use:

```
sudo dpkg --get-selections | grep <search-term>
```

If the packages are not installed, they can be installed with:

```
sudo apt update
sudo apt install make binutils gcc build-essential linux-libc-dev pahole fakeroot
flex binutils-dev
sudo apt install --reinstall gcc-13
```

After performing package installations, it may be necessary and helpful to reboot the VM.

1. Download kernel module starter code.

To begin, download the tar.gz archive containing a basic Hello World Linux kernel module.

First, create a directory for this assignment, and then download the tar gz:

Kernel modules will FAIL TO COMPILE if source code is placed under paths with spaces.
In general spaces in filenames is a Windows/MAC convention and is uncommon on Linux.
Please do not use spaces in directory or filenames.

```
mkdir a3
cd a3
```

```
wget http://faculty.washington.edu/wlloyd/courses/tcss422/assignments/hello\_module.tar.gz
tar xzf hello_module.tar.gz
cd hello_module
ls -l
```

The hello_module contains “helloModule.c”, which includes source code for the most basic Linux kernel module. Inspect the code. The module consists of an initialization method and a cleanup method. The initialization method is triggered automatically as an “event” when the module is loaded. The cleanup method is triggered when the kernel module is unloaded. Kernel modules are loaded dynamically into the operating system. They are not run as user programs on the command line. Consequently, they don’t have traditional I/O (e.g. stdin, stdout).

Note: Required #include files have already been included in hello_module.c for this assignment.

2. Compile, install, and test a Linux Kernel Module

Before compiling the kernel module, to suppress the BTF compiler warning, run this command:

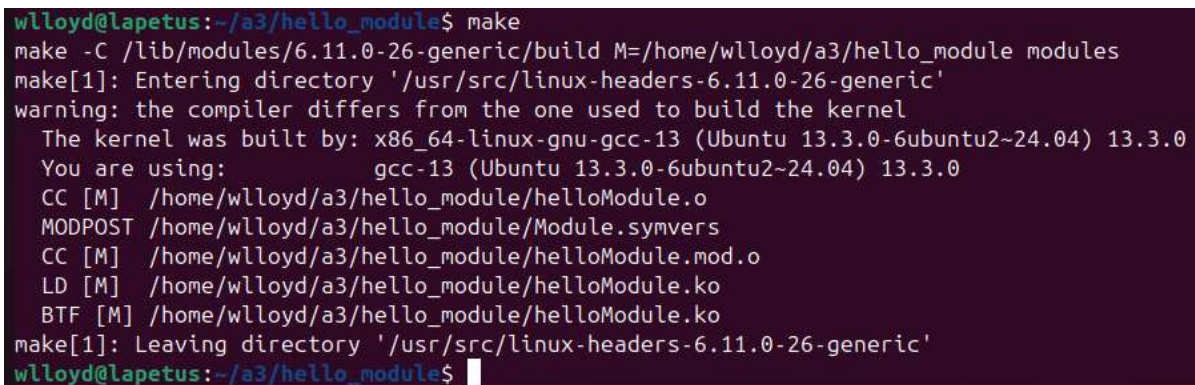
```
sudo ln -s /sys/kernel/btf/vmlinux /usr/lib/modules/$(uname -r)/build/
```

BTF is BPF Type Format, a metadata format that encodes debug information for eBPF programs and the Linux kernel, including kernel modules. When building the kernel module, Ubuntu 24.04 appears to be missing a path to /sys/kernel/btf/vmlinux which causes a warning message and BTF generation to be skipped.

Try compiling this module, and inspecting the resulting files:

```
make
```

You may receive the compiler warning as below:



```
wlloyd@lapetus:~/a3/hello_module$ make
make -C /lib/modules/6.11.0-26-generic/build M=/home/wlloyd/a3/hello_module modules
make[1]: Entering directory '/usr/src/linux-headers-6.11.0-26-generic'
warning: the compiler differs from the one used to build the kernel
The kernel was built by: x86_64-linux-gnu-gcc-13 (Ubuntu 13.3.0-6ubuntu2~24.04) 13.3.0
You are using: gcc-13 (Ubuntu 13.3.0-6ubuntu2~24.04) 13.3.0
CC [M] /home/wlloyd/a3/hello_module/helloModule.o
MODPOST /home/wlloyd/a3/hello_module/Module.symvers
CC [M] /home/wlloyd/a3/hello_module/helloModule.mod.o
LD [M] /home/wlloyd/a3/hello_module/helloModule.ko
BTF [M] /home/wlloyd/a3/hello_module/helloModule.ko
make[1]: Leaving directory '/usr/src/linux-headers-6.11.0-26-generic'
wlloyd@lapetus:~/a3/hello_module$
```

The compiler warning can be ignored as long as the kernel module compiles. Check if the module compiled:

```
ls -alt
```

If gcc-13 is NOT installed, the kernel module may fail to compile.

Check if gcc-13 is installed:

```
$ whereis gcc-13
```

```
gcc-13: /usr/bin/gcc-13 /usr/share/man/man1/gcc-13.1.gzs
```

If gcc-13 is not installed, install it as described on page 1.

The newly compiled kernel module should appear at the top of the directory output. Compilation produces a kernel object (ko) file. Load this file:

```
sudo insmod helloModule.ko
```

The “insmod” command must be run as superuser. This “inserts” a module into the running kernel. When we load the module the init() method is automatically called. Any “printk” statements are sent to the kernel log file. To view what has been written to the kernel log file, check the last 10 lines of /var/log//syslog using:

```
tail -n 10 /var/log/syslog
```

The module can be unloaded using the “rmmod” command. It is not necessary to include the “ko” file extension:

```
sudo rmmod helloModule
```

Check the kernel log after unloading the module:

```
tail -n 10 /var/log/syslog
```

Separate log messages are written by the init() and cleanup() methods.

For this tutorial, we will iteratively make changes to the helloModule kernel module. For each iteration we will “rmmod” the previous version, “insmod” the new version”, and check printk output to /var/log/syslog using the tail command.

Messages regarding kernel module verification can be ignored as this is a user provided test module.

3. Count the number of running Linux tasks

Only privileged kernel code can read kernel data structures.

Let’s access the task_struct data structure to inspect running processes and threads on the computer. Let’s add code to the init() method that is run when the module is loaded.

The task_struct data structure requires the <linux/sched/signal.h> include statement which defines an important helper function called for_each_process(). This include statement is **already** in the starter code.

```
#include<linux/sched/signal.h>
```

Now add source code for the function:

Under “Helper Routine Function Prototypes”, add the function prototype:

```
int proc_count(void);
```

Then under “Helper Routine implementation”, add the `proc_count()` function implementation:

```
int proc_count(void)
{
    int i=0;
    struct task_struct *thechild;
    for_each_process(thechild)
        i++;
    return i;
}
```

Next, call this function by adding a `printk()` statement to `init()`:

```
printk(KERN_INFO "There are %d running processes.\n", proc_count());
```

In this assignment, we leverage two Linux kernel data structures:

```
struct task_struct
defined in /usr/src/linux-headers-$(uname -r)/include/linux/sched.h
starting at line 758 (value for 6.11 Ubuntu 24.04 kernel, exact location may vary)

struct mm_struct
defined in /usr/src/linux-headers-$(uname -r)/include/linux/mm_types.h
starting at line 809 (value for 6.11 Ubuntu 24.04 kernel, exact location may vary)
```

Before recompiling, first check to ensure that these files are on your system with “ls”:

```
ls -l /usr/src/linux-headers-$(uname -r)/include/linux/sched.h
ls -l /usr/src/linux-headers-$(uname -r)/include/linux/mm_types.h
```

If you do not have these files, check if they are installed with:

```
sudo dpkg -l | grep linux-headers-$(uname -r)
```

If the result is blank, then try installing the required headers:

```
sudo apt install linux-headers-$(uname -r) linux-headers-$(uname -r)-generic
```

Then test if they are available:

```
sudo dpkg -l | grep linux-headers-$(uname -r)
```

If there are still problems obtaining the linux headers, please [contact the instructor](#).

Recompile and reload the kernel module and check output to the syslog.

<Q1> - How many processes do you see?

4. Inspecting Linux Task Codesize

Next, implement the function below to scan Linux processes and determine the code size for each of them.

Important Note about COPY-AND-PASTE:

If copying and pasting code from a PDF file, please note that some characters may not be copied correctly. It may be necessary to correct quote marks, subtract operators, etc. **The DOCX file is better for copy/paste.**

Now add another helper function. Be sure to add the function prototype as before.

```
unsigned long code_memory(void)
{
    unsigned long totalmem=0;
    unsigned long thismem;
    struct task_struct *thechild;
    for_each_process(thechild)
    {
        printk(KERN_INFO "pid %d", thechild->pid);
        thismem=thechild->mm->end_code - thechild->mm->start_code;
        printk(KERN_INFO " has codesize %lu.\n", thismem);
        totalmem=totalmem+thismem;
    }
    return totalmem;
}
```

Now, call this function by adding a printk() statement to init():

```
printk(KERN_INFO "Running processes have %lu bytes of code.\n", code_memory());
```

Recompile and reload the kernel module and check output to the syslog.

<Q2> – What happens when this kernel module is loaded?

Check the kernel log file.

```
tail -n 60 /var/log/syslog
```

Scroll up to view the full content of the log.

Look for an error message that says “BUG: unable to handle kernel”...

Examine the description of the error carefully to learn what has caused the issue.

>> Now, reboot and restore the Ubuntu virtual machine before continuing. <<

To proceed the bug needs to be corrected.

Not every process can be indexed with the code_memory() function.

Processes must be tested before running the code.

Line 9 of the function dereferences **two pointers**:

```
    thismem=thechild->mm->end_code - thechild->mm->start_code;
```

When the first pointer is not valid, it is not possible to dereference the second pointer.

When the second pointer is not valid, it is not possible to read any members of the mm (memory map).

The variable “thechild” is a task_struct pointer.

Let's inspect what member "mm" is by looking at sched.h in the task_struct type definition. See the table above on page 4 for the location of the **sched.h include file**. Skip to approximately **line 898** in the file.

Line 898 assumes Ubuntu 24.04 with Linux kernel 6.11.0-26.

If using a different version of Ubuntu, this line number could vary.

You'll be looking first for the location where struct task_struct is defined – and then the "mm" pointer within task_struct.

The definition of task_struct will look like:

```
struct task_struct {
#ifdef CONFIG_THREAD_INFO_IN_TASK
    /*
     * For reasons of header soup (see current_thread_info()), this
     * must be the first element of task_struct.
     */
    struct thread_info thread_info;
#endif
    volatile long state; /* -1 unrunnable, 0 runnable, >0 stopped */
    void *stack;
```

To correct the bug, we must first check that mm is valid before using it.

<Q3> – What test should be made before dereferencing a member that is pointed to in C ?

Correct the issue in Q3, and then continue.

The correction should allow the function to run where it will print the codesize of every process to the /var/log/syslog file.

***** If you experience trouble correcting the error, contact the instructor for a hint. *****

Next, for simplicity **combine** the two printk() statements in the function.

Remove the first printk() statement, and replace the second printk() with the following:

```
printk(KERN_INFO "pid %d has codesize %lu.\n", thechild->pid, thismem);
```

Recompile and reload the kernel module.

Test the output by inspecting the tail of the log file:

```
tail -n 250 /var/log/syslog
```

It should be possible to see processes with and without code:

```
2025-05-26T23:35:35.868235-07:00 lapetus kernel: pid 3226 has codesize 29821.
2025-05-26T23:35:35.868235-07:00 lapetus kernel: pid 3260 has codesize 148245.
2025-05-26T23:35:35.868235-07:00 lapetus kernel: pid 3277 has codesize 6481.
2025-05-26T23:35:35.868236-07:00 lapetus kernel: pid 3302 has codesize 6481.
2025-05-26T23:35:35.868237-07:00 lapetus kernel: pid 3314 has codesize 6481.
2025-05-26T23:35:35.868237-07:00 lapetus kernel: pid 3344 has codesize 30833.
2025-05-26T23:35:35.868237-07:00 lapetus kernel: pid 3347 has codesize 292893.
2025-05-26T23:35:35.868237-07:00 lapetus kernel: pid 3348 has codesize 449405.
2025-05-26T23:35:35.868238-07:00 lapetus kernel: pid 3373 has codesize 242133.
2025-05-26T23:35:35.868238-07:00 lapetus kernel: pid 3416 has codesize 292093.
2025-05-26T23:35:35.868238-07:00 lapetus kernel: pid 3447 has codesize 73873.
2025-05-26T23:35:35.868238-07:00 lapetus kernel: pid 3457 has codesize 11329.
2025-05-26T23:35:35.868239-07:00 lapetus kernel: pid 3463 has codesize 132873.
2025-05-26T23:35:35.868239-07:00 lapetus kernel: pid 3587 has codesize 152621.
2025-05-26T23:35:35.868239-07:00 lapetus kernel: pid 3594 has codesize 975661.
2025-05-26T23:35:35.868240-07:00 lapetus kernel: pid 3676 has codesize 37217.
2025-05-26T23:35:35.868240-07:00 lapetus kernel: pid 3771 has codesize 37217.
2025-05-26T23:35:35.868240-07:00 lapetus kernel: pid 4083 has codesize 37217.
2025-05-26T23:35:35.868240-07:00 lapetus kernel: pid 4422 has codesize 195537.
2025-05-26T23:35:35.868241-07:00 lapetus kernel: pid 4423 has codesize 195537.
2025-05-26T23:35:35.868241-07:00 lapetus kernel: pid 4424 has codesize 115077.
2025-05-26T23:35:35.868241-07:00 lapetus kernel: Running processes have 17124296
8 bytes of code.
```

<Q4> – What is the codesize of process #1? This is the Linux kernel’s init process, the parent of all processes.
Write down the codesize of process #1.

Now, in the Linux include files, search for the “mm” member of this data structure.

In the task_struct definition, the mm member may be defined around **line 898** of the file.

<Q5> – What is the data type for the mm member of task_struct data structure?
Provide name exactly as it appears in the header file. (case sensitive)

5. Inspecting Linux Task Datasize and Heapsize

Next, look up the structure for task_struct’s “mm” member.
This data type (struct) is defined in the **mm_types.h include file**.
Check the table above on page 4 for the location of this file.

This struct, defined in mm_types, spans many lines of code.
Look for the two members (fields) in the struct that we’ve used in the code above:

```
field: start_code  
field: end_code
```

Now copy the code_memory() function and create two new functions.
Name one function: “**data_memory()**” and the other “**heap_memory()**”.

Replace the use of start_code and end_code in these functions.

In data_memory, use:

```
field: start_data  
field: end_data
```

In heap_memory, use:

```
field: start_brk  
field: brk
```

In the init() function modify the printk call to call all three memory functions:

```
printk(KERN_INFO "Running processes have %lu bytes of code, %lu bytes of data, and  
%lu bytes of heap.\n", code_memory(), data_memory(), heap_memory());
```

Rebuild and reinstall the kernel module, and then check the output to /var/log/syslog.
There will be a lot of output lines. Try adjusting the number of lines returned by the tail command to capture data only from your latest version of the kernel module. There should be data for ‘heapsize’, ‘datasize’, and ‘codesize’ and then summary statistics. Adjust the number of lines returned so you do not see any output from the prior runs. This will be roughly ~ 3x the number of processes plus 10.

For example:

```
$tail -n 650 /var/log/syslog
```

<Q6> - For PID #1, which segment type utilizes the smallest amount of memory space?

- A. Heap
- B. Data
- C. Code
- D. Stack
- E. None of the above

<Q7> - For all processes combined, which segment type utilizes the largest amount of memory space?

- A. Heap
- B. Data
- C. Code
- D. Stack
- E. None of the above

<Q8> - For all processes combined, which segment type utilizes the smallest amount of memory space?

- A. Heap
- B. Data
- C. Code
- D. Stack
- E. None of the above

<Q9> - Please copy and paste the complete output from your final Linux kernel module from `/var/log/syslog` into the box below for Question 9. Be sure that the output captures 'heapsize', 'datasize', and 'codesize' of every process, plus the summary statistics. Do not include output from prior testing of the kernel module.

<Q10> - (required) Upload the entire unmodified `/var/log/syslog` file to Canvas. A web browser such as Firefox can be used on your VM to interact directly with Canvas. Alternatively, try using a web browser to email yourself the `syslog` file as an attachment, and the submit it to Canvas.

To complete this assignment, to obtain credit, log in to Canvas, and complete the **Assignment 3 Quiz**.

Generative-AI policy:

Only Generative-AI output produced by Google Search is permitted for this assignment. Google Search combines several LLMs including Gemini to produce intelligent search results. It is not necessary to document the use of LLMs in this assignment. Since this assignment largely requires hands-on step-by-step effort to build kernel modules, it is unlikely that an LLM will be very helpful.

| Version | Date | Change |
|---------|----------|------------------|
| 0.1 | 3/5/2026 | Original Version |

Please help improve this tutorial by reporting any errors or discrepancies found.