

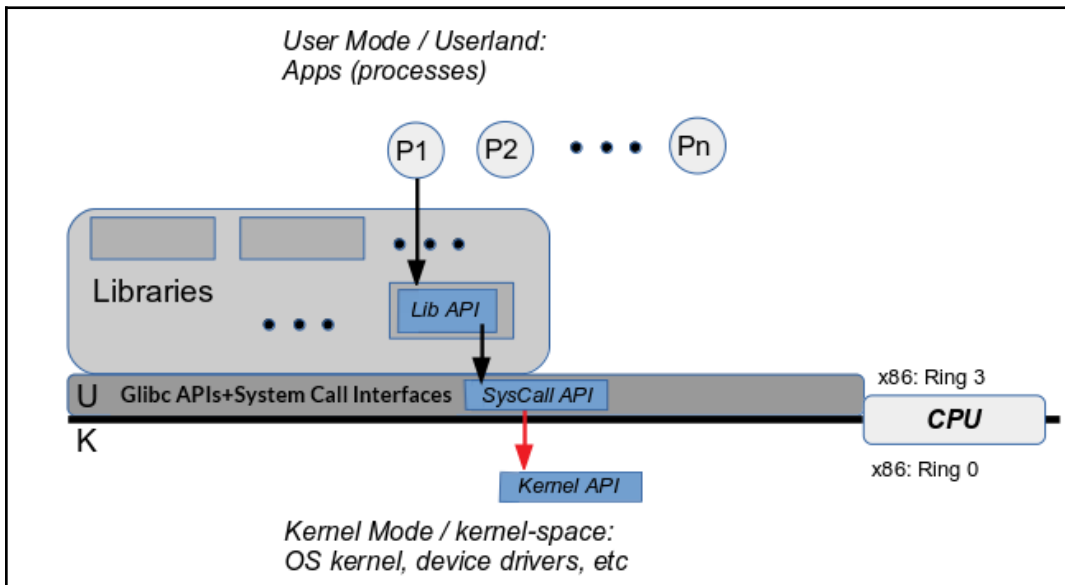
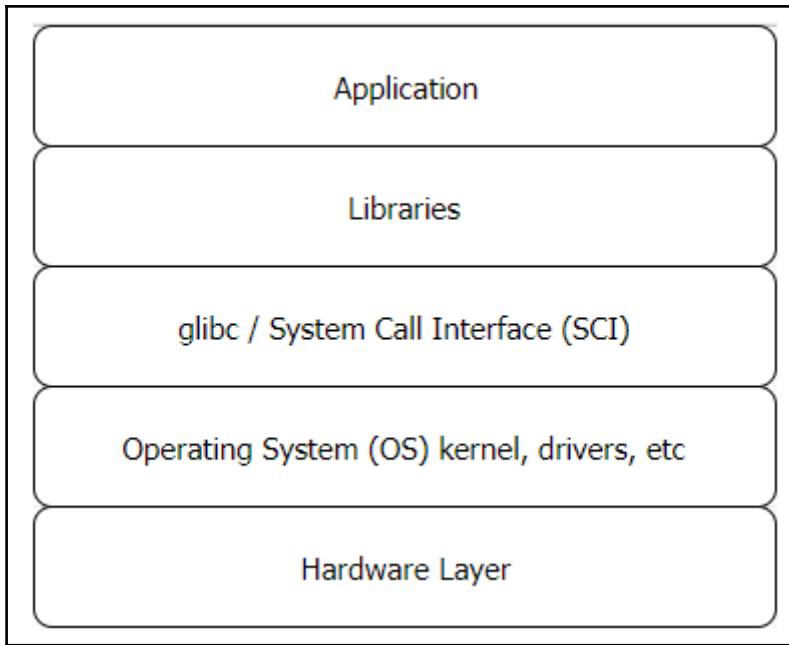
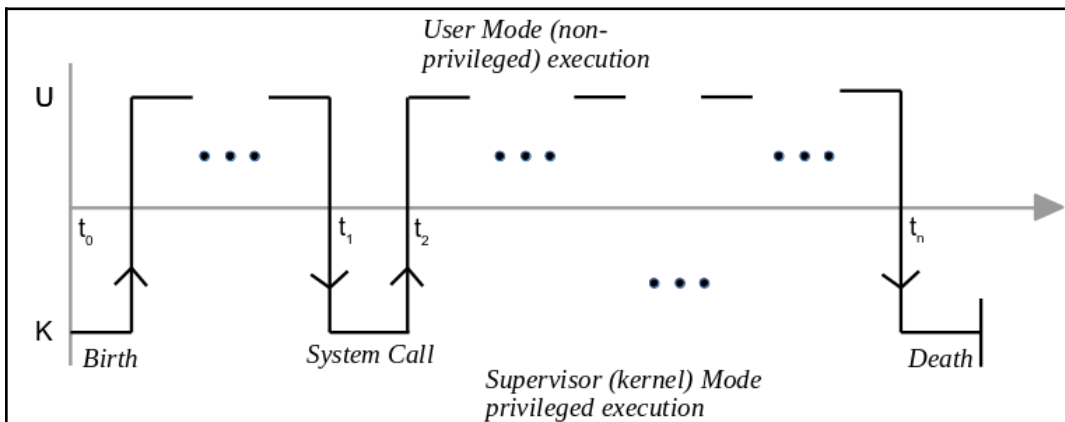
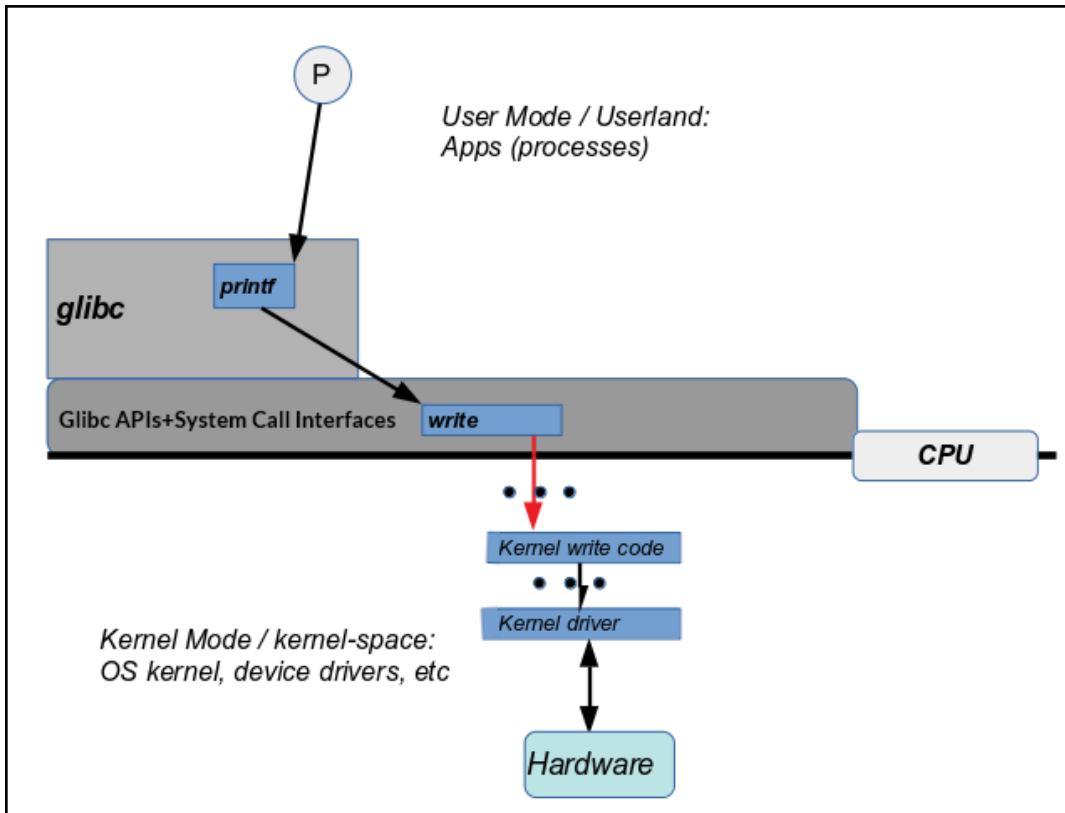


Chapter 1: Linux System Architecture






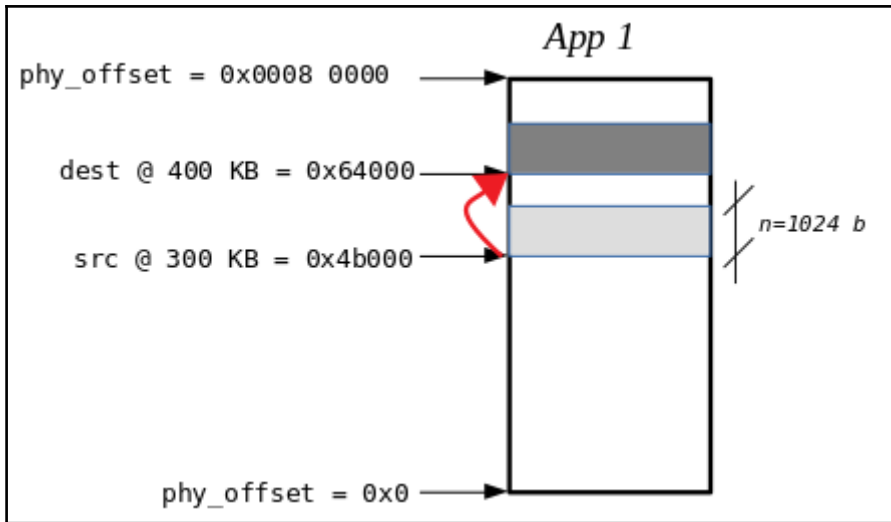






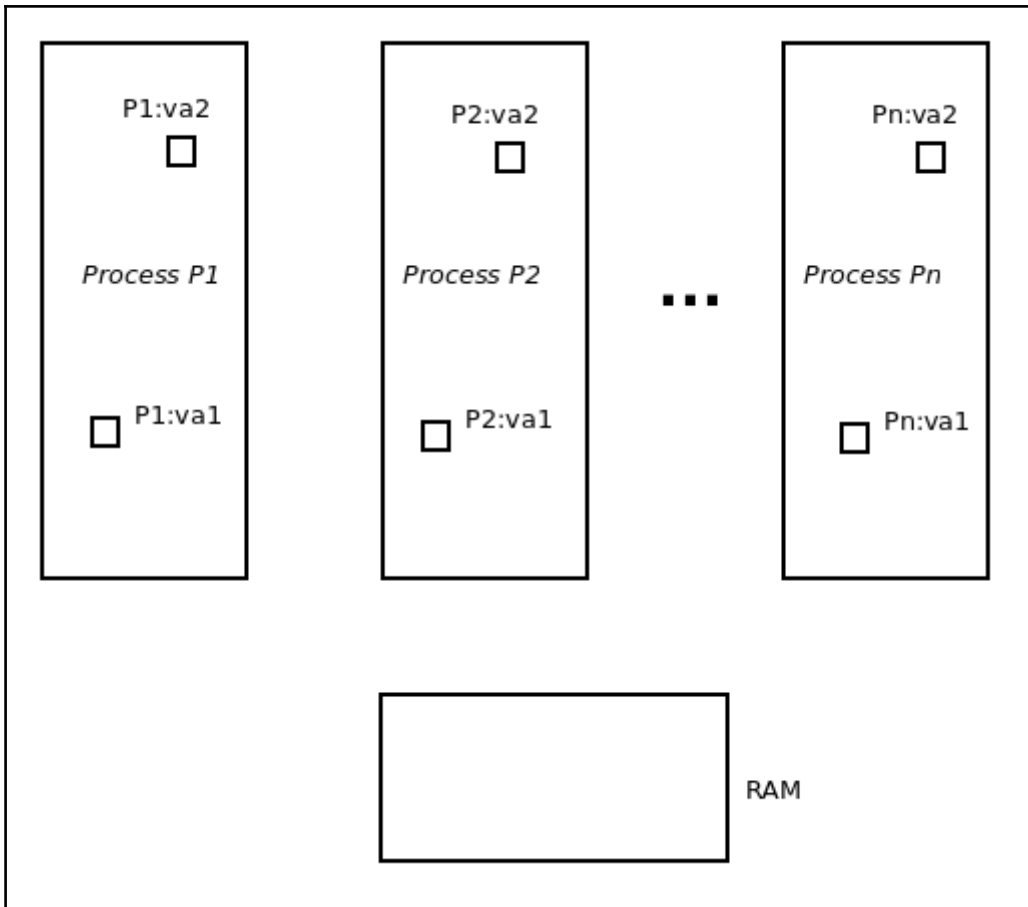


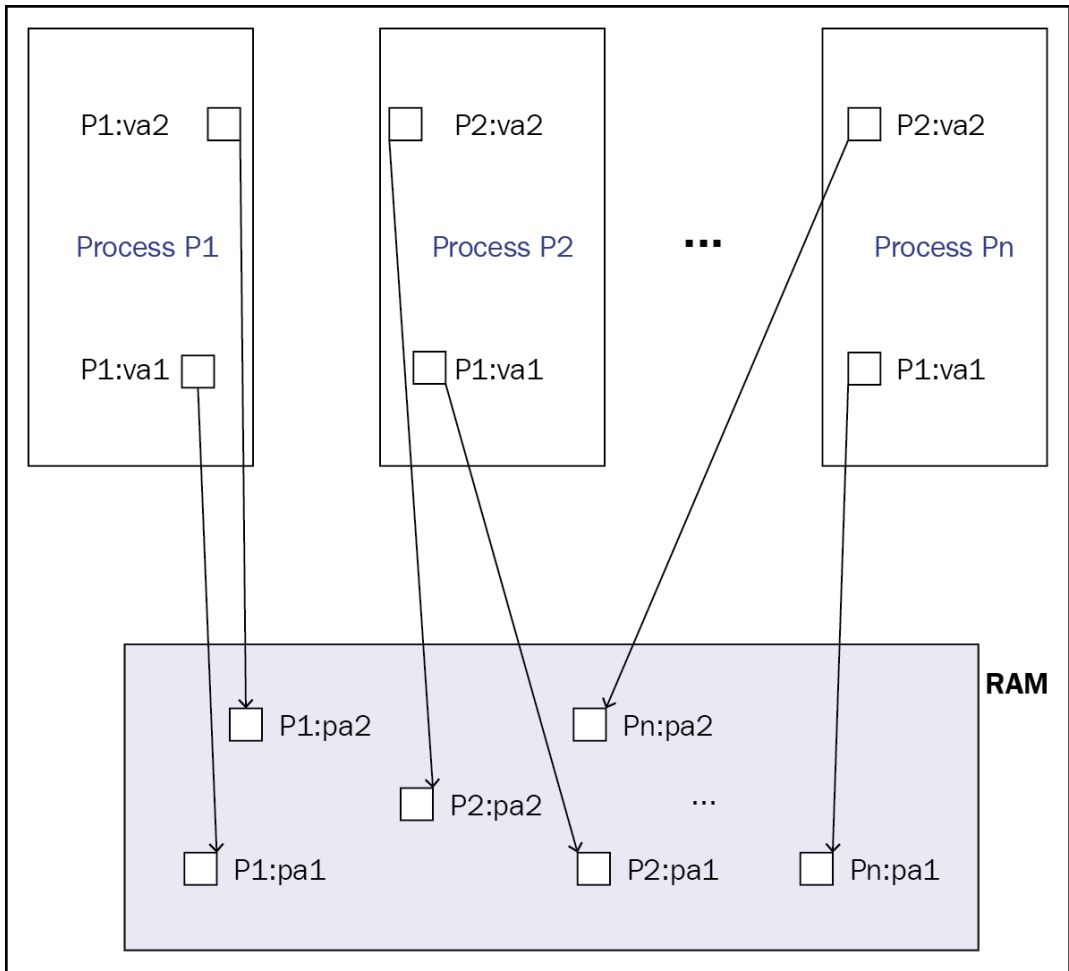
Chapter 2: Virtual Memory

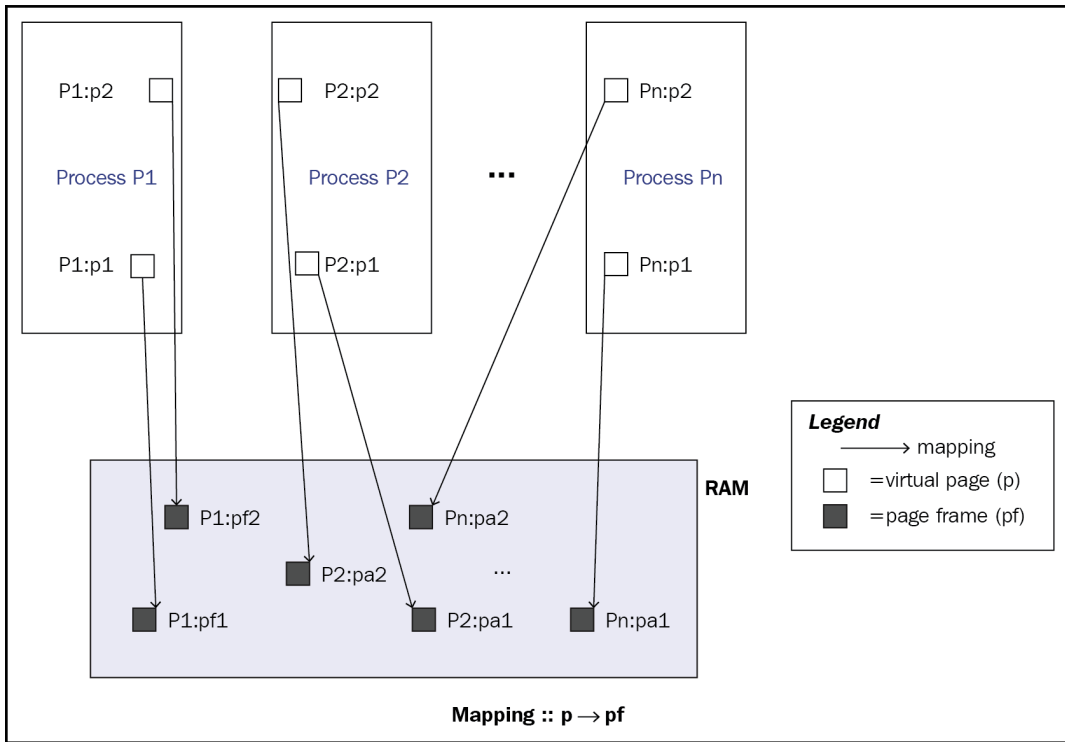





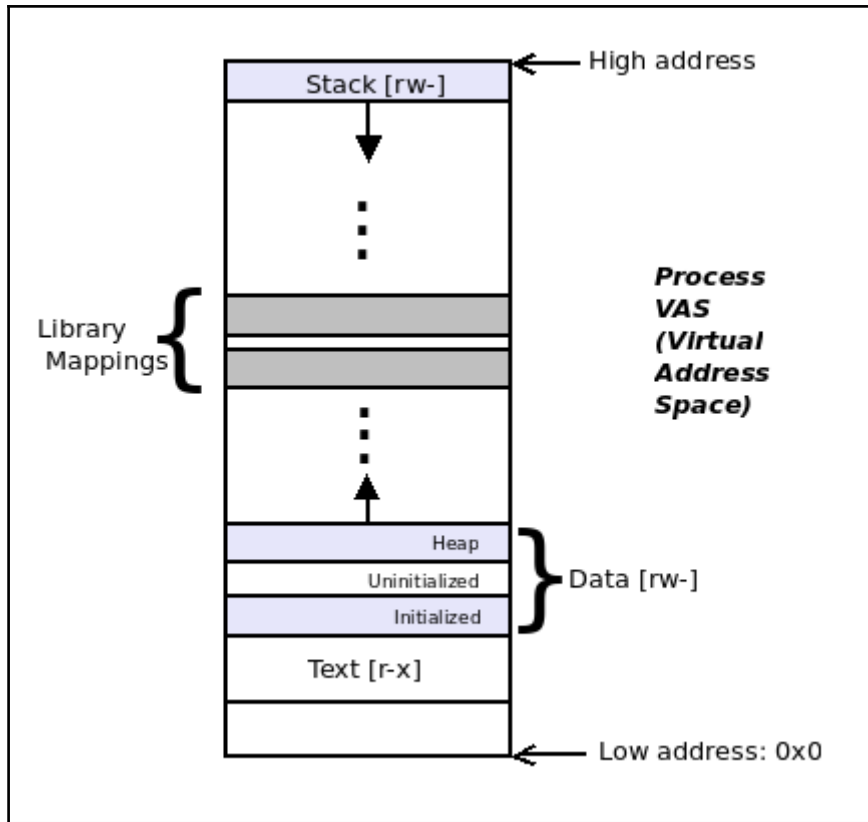


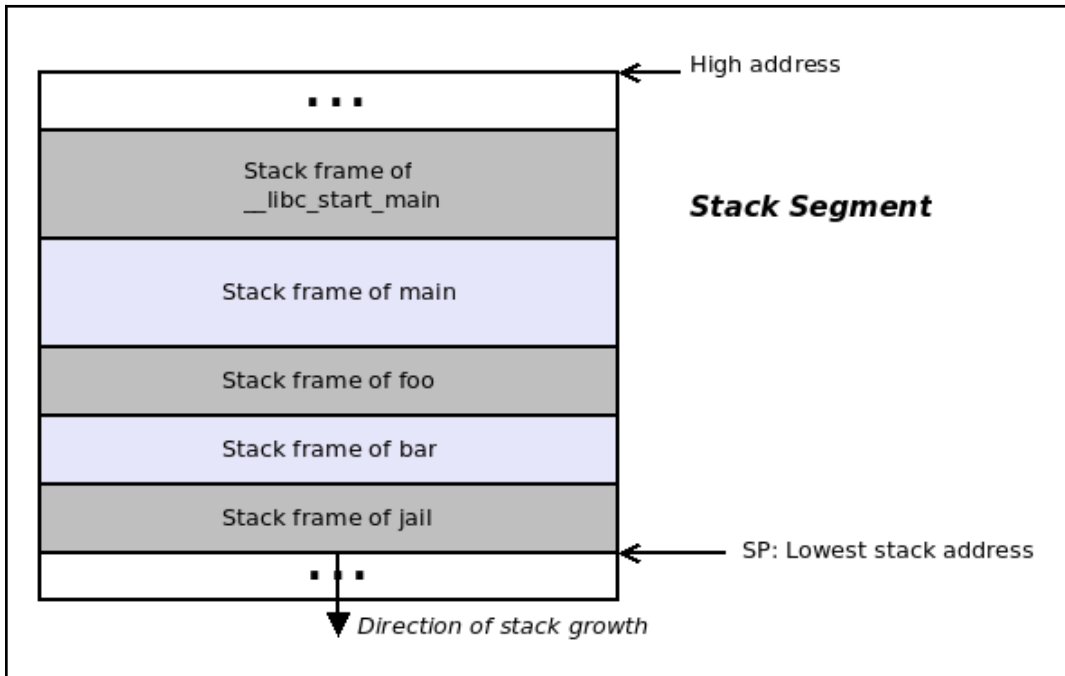
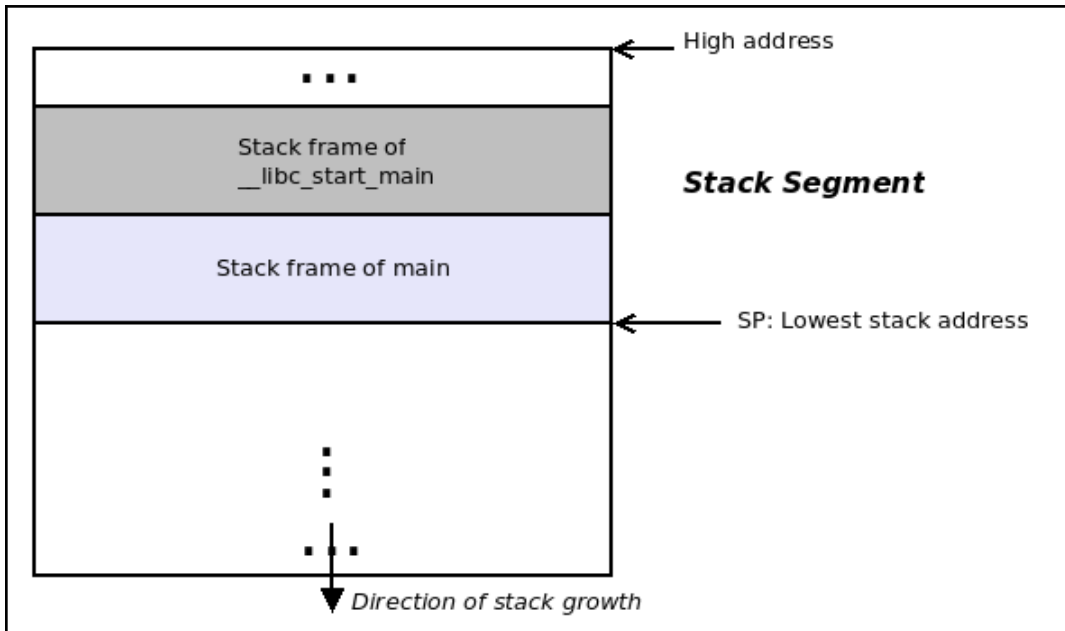


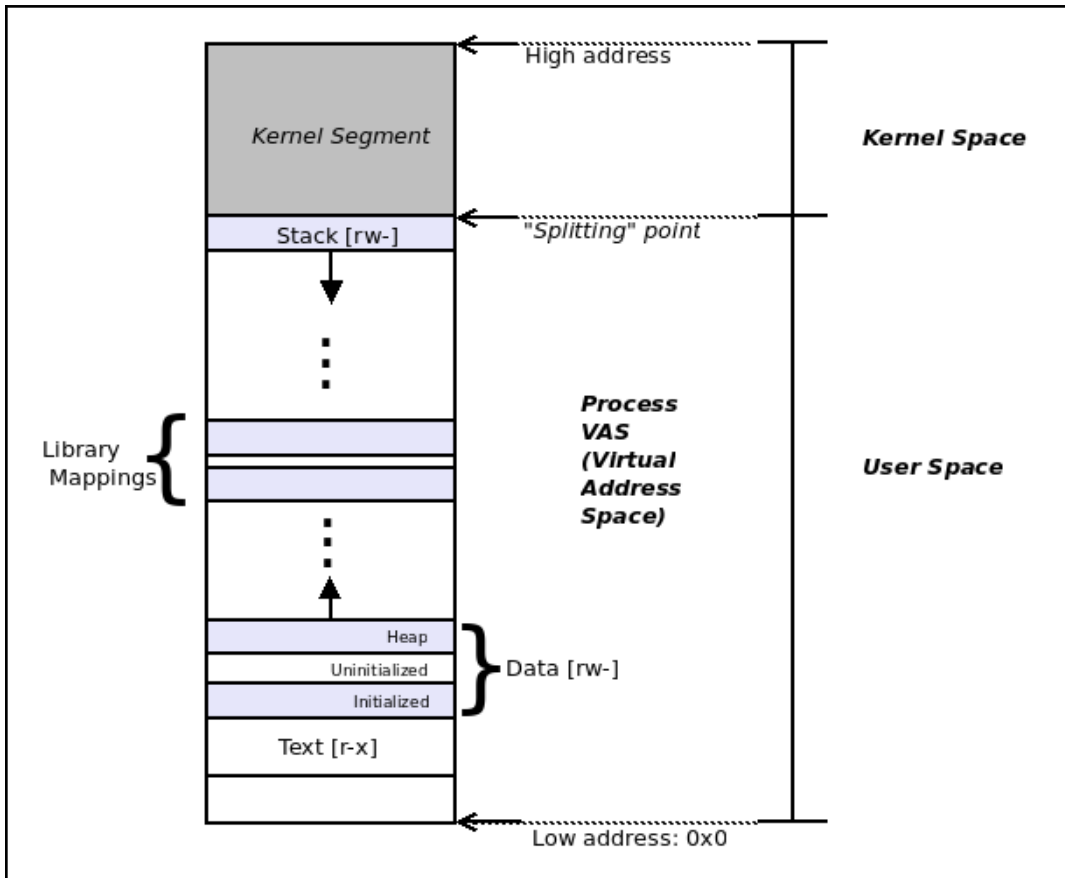


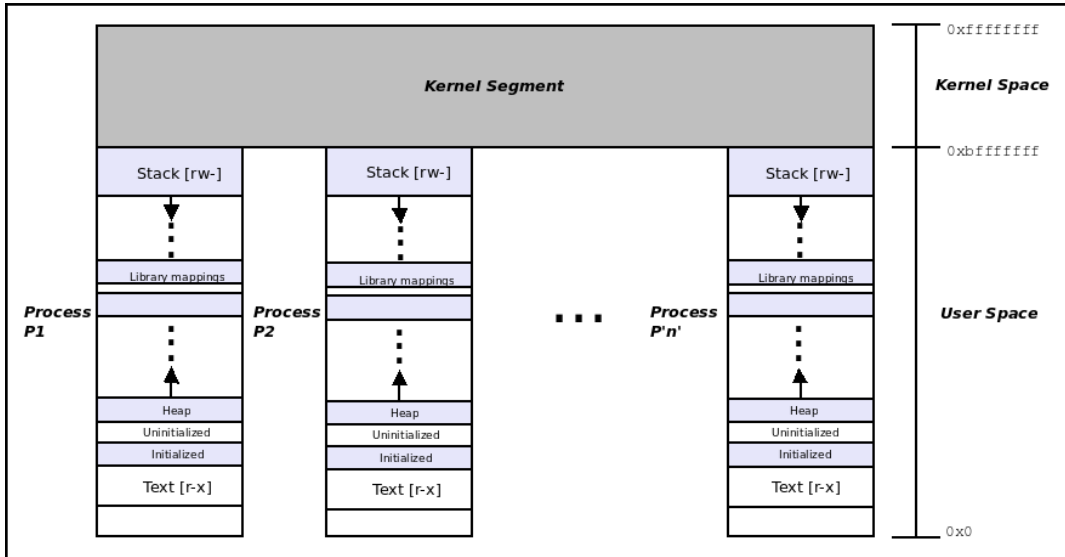






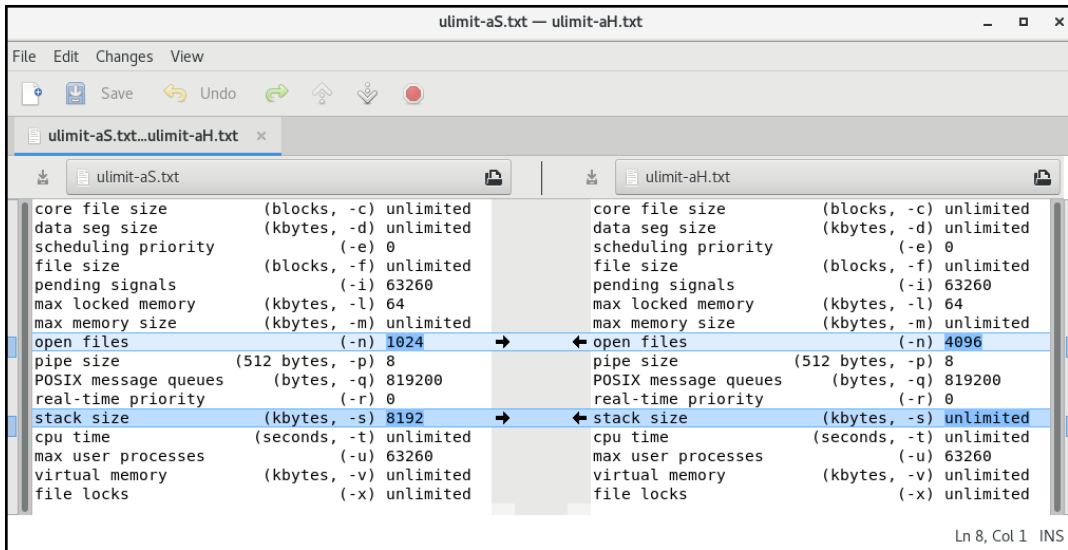






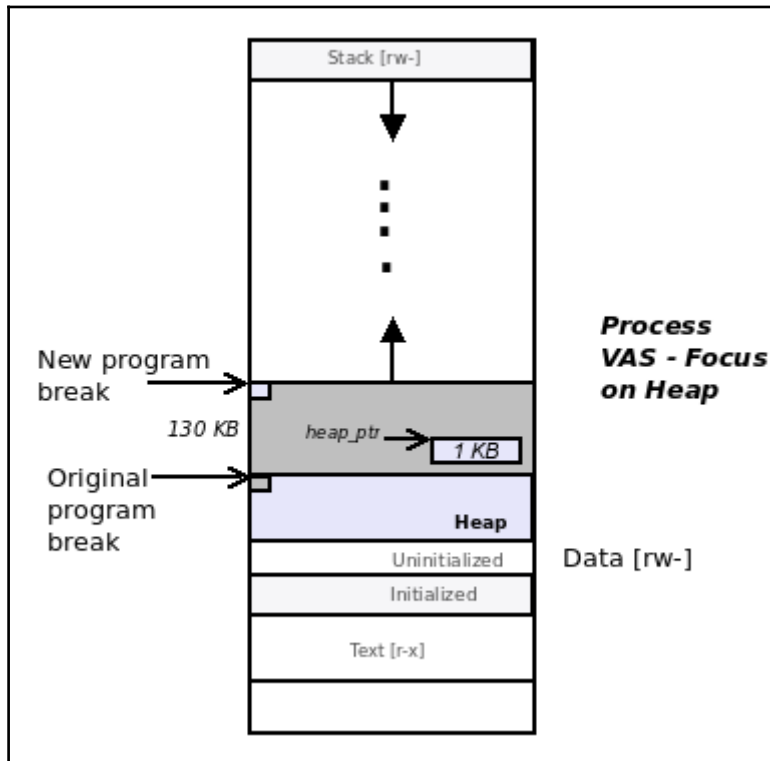
Chapter 3: Resource Limits

```
$ ulimit -f
unlimited
$ ulimit -f 2
$ ulimit -f
2
$ dd if=/dev/urandom of=tst count=2048 bs=1
2048+0 records in
2048+0 records out
2048 bytes (2.0 kB, 2.0 KiB) copied, 0.00688134 s, 298 kB/s
$ dd if=/dev/urandom of=tst count=2049 bs=1
File size limit exceeded (core dumped)
$ █
```



RESOURCE	DESCRIPTION	SOFT	HARD	UNITS
AS	address space limit	unlimited	unlimited	bytes
CORE	max core file size	0	unlimited	blocks
CPU	CPU time	unlimited	unlimited	seconds
DATA	max data size	unlimited	unlimited	bytes
FSIZE	max file size	2048	2048	blocks
LOCKS	max number of file locks held	unlimited	unlimited	
MEMLOCK	max locked-in-memory address space	65536	65536	bytes
MSGQUEUE	max bytes in POSIX mqueues	819200	819200	bytes
NICE	max nice prio allowed to raise	0	0	
NOFILE	max number of open files	2000	2000	
NPROC	max number of processes	7741	7741	
RSS	max resident set size	unlimited	unlimited	pages
RTPRIO	max real-time priority	0	0	
RTTIME	timeout for real-time tasks	unlimited	unlimited	microsecs
SIGPENDING	max number of pending signals	7741	7741	
STACK	max stack size	8388608	unlimited	bytes

Chapter 4: Dynamic Memory Allocation



```

22875 3) memprot-29591 | d...+ 10.204 us | }
22876 3) memprot-29591 | d... | do syscall 64() {
22877 3) memprot-29591 | .... |   Sys_mprotect() {
22878 3) memprot-29591 | .... |     do_mprotect_pkey() {
22879 3) memprot-29591 | .... |       down_write_killable() {
22880 3) memprot-29591 | .... |         _cond_resched() {
22881 3) memprot-29591 | .... 0.034 us |           rcu_all_qs();
22882 3) memprot-29591 | .... 0.321 us |         }
22883 3) memprot-29591 | .... 0.616 us |       }
22884 3) memprot-29591 | .... |     }
22885 3) memprot-29591 | .... 0.107 us |   find_vma() {
22886 3) memprot-29591 | .... 0.116 us |     vmacache_find();
22887 3) memprot-29591 | .... 1.227 us |     vmacache_update();
22888 3) memprot-29591 | .... |   }
22889 3) memprot-29591 | .... | security_file_mprotect() {
22890 3) memprot-29591 | .... |   selinux_file_mprotect() {
22891 3) memprot-29591 | .... 0.047 us |     avc_has_perm() {
22892 3) memprot-29591 | .... |     avc_denied();
22893 3) memprot-29591 | .... |     slow_avc_audit() {
                                     |     common_lsm_audit() {

```

Chapter 5: Linux Memory Issues

There are **91** matching records.
Displaying matches **1** through **20**.

1 2 3 4 5 > >>

Vuln ID 	Summary 	CVSS Severity 
CVE-2017-18120	A double-free bug in the read_gif function in gifread.c in gifsicle 1.90 allows a remote attacker to cause a denial-of-service attack or unspecified other impact via a maliciously crafted file, because last_name is mishandled, a different vulnerability than CVE-2017-1000421. Published: February 02, 2018; 04:29:00 AM -05:00	(not available)
CVE-2018-0101	A vulnerability in the Secure Sockets Layer (SSL) VPN functionality of the Cisco Adaptive Security Appliance (ASA) Software could allow an unauthenticated, remote attacker to cause a reload of the affected system or to remotely execute code. The vulnerability is due to an attempt to double free a region of memory when the webvpn feature is enabled on the Cisco ASA device. An attacker could exploit this vulnerability by sending	(not available)

Chapter 6: Debugging Tools for Memory Issues

```
$ valgrind ./membugs_dbg 8
==24337== Memcheck, a memory error detector
==24337== Copyright (C) 2002-2017, and GNU GPL'd, by Julian Seward et al.
==24337== Using Valgrind-3.13.0 and LibVEX; rerun with -h for copyright info
==24337== Command: ./membugs_dbg 8
==24337==
uaf():165: arr = 0x521f040:aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
==24337== Invalid write of size 1
==24337==    at 0x4C32E7F: strncpy (vg_replace_strmem.c:549)
==24337==    by 0x400FB5: uaf (membugs.c:173)
==24337==    by 0x4014A6: process_args (membugs.c:348)
==24337==    by 0x401574: main (membugs.c:379)
==24337== Address 0x521f040 is 0 bytes inside a block of size 512 free'd
==24337==    at 0x4C30D18: free (vg_replace_malloc.c:530)
==24337==    by 0x400F93: uaf (membugs.c:171)
==24337==    by 0x4014A6: process_args (membugs.c:348)
==24337==    by 0x401574: main (membugs.c:379)
==24337== Block was alloc'd at
==24337==    at 0x4C2FB6B: malloc (vg_replace_malloc.c:299)
==24337==    by 0x400EC5: uaf (membugs.c:159)
==24337==    by 0x4014A6: process_args (membugs.c:348)
==24337==    by 0x401574: main (membugs.c:379)
==24337==
==24337== Invalid read of size 1
==24337==    at 0x4C32B2B: strlen (vg_replace_strmem.c:425)
==24337==    by 0x4E8FFB2: vfprintf (in /usr/lib64/libc-2.26.so)
==24337==    by 0x4E99255: printf (in /usr/lib64/libc-2.26.so)
==24337==    by 0x400FE2: uaf (membugs.c:174)
==24337==    by 0x4014A6: process_args (membugs.c:348)
==24337==    by 0x401574: main (membugs.c:379)
```

```

$ ./membugs_dbg_asan 2
=====
==4125==ERROR: AddressSanitizer: stack-buffer-overflow on address 0x7ffe9f8d8174 at pc 0x00000051271d bp 0x7ffe9f8d8130 sp 0x7ffe9f8d8128
WRITE of size 4 at 0x7ffe9f8d8174 thread T0
#0 0x51271c (/home/seawolf/tmp/membugs_dbg_asan+0x51271c)
#1 0x51244e (/home/seawolf/tmp/membugs_dbg_asan+0x51244e)
#2 0x512291 (/home/seawolf/tmp/membugs_dbg_asan+0x512291)
#3 0x7fe3408ccb96 (/lib/x86_64-linux-gnu/libc.so.6+0x21b96)
#4 0x419ea9 (/home/seawolf/tmp/membugs_dbg_asan+0x419ea9)

Address 0x7ffe9f8d8174 is located in stack of thread T0 at offset 52 in frame
#0 0x5125ef (/home/seawolf/tmp/membugs_dbg_asan+0x5125ef)

This frame has 1 object(s):
[32, 52) 'arr' (line 265) <== Memory access at offset 52 overflows this variable
HINT: this may be a false positive if your program uses some custom stack unwind mechanism or swapcontext
(longjmp and C++ exceptions *are* supported)
SUMMARY: AddressSanitizer: stack-buffer-overflow (/home/seawolf/tmp/membugs_dbg_asan+0x51271c)
Shadow bytes around the buggy address:
 0x100053f12fd0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 0x100053f12fe0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 0x100053f12ff0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 0x100053f13000: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 0x100053f13010: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
=>0x100053f13020: 00 00 00 00 00 00 00 00 f1 f1 f1 00 00[04]f3
 0x100053f13030: f3 f3 f3 f3 00 00 00 00 00 00 00 00 00 00 00
 0x100053f13040: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 0x100053f13050: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 0x100053f13060: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 0x100053f13070: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Shadow byte legend (one shadow byte represents 8 application bytes):
Addressable:          00

```

```

$ ./membugs_dbg_asan 13

## Leakage test: case 3: "lib" API: runtime cond = 0
mypath = /usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin:/usr/games:/usr/local/games:/snap/bin:/sbin:/usr/sbin:/usr/local/sbin:/home/seawolf/kaiwanTECH/usefulsnips:/usr/lib/llvm-6.0/bin/

## Leakage test: case 3: "lib" API: runtime cond = 1
mypath = /usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin:/usr/games:/usr/local/games:/snap/bin:/sbin:/usr/sbin:/usr/local/sbin:/home/seawolf/kaiwanTECH/usefulsnips:/usr/lib/llvm-6.0/bin/

=====
==26220==ERROR: LeakSanitizer: detected memory leaks

Direct leak of 4096 byte(s) in 1 object(s) allocated from:
#0 0x4d9d60 in malloc (/home/seawolf/tmp/membugs_dbg_asan+0x4d9d60)
#1 0x513eca in silly_getpath /home/seawolf/kaiwanTECH/book_src/ch5/membugs.c:37:9
#2 0x513d0d in leakage_case3 /home/seawolf/kaiwanTECH/book_src/ch5/membugs.c:55:2
#3 0x512522 in process_args /home/seawolf/kaiwanTECH/book_src/ch5/membugs.c:363:4
#4 0x512291 in main /home/seawolf/kaiwanTECH/book_src/ch5/membugs.c:375:2
#5 0x7f5f23247b96 in __libc_start_main /build/glibc-0TSEL5/glibc-2.27/csu/../csu/libc-start.c:310

SUMMARY: AddressSanitizer: 4096 byte(s) leaked in 1 allocation(s).
$ █

```

Chapter 7: Process Credentials

Access Category	Owner (U)			Group (G)			Others (O)		
Permission Bits	r	w	x	r	w	x	r	w	x
Example (myfile)	1	1	0	1	1	0	1	0	0

```
$ ls -l $(which passwd)
-rwsr-xr-x 1 root root 54224 Aug 21 05:26 /usr/bin/passwd*
$ █
```

system(3) - Linux manual page - Mozilla Firefox (Private Browsing)

system(3) - Linux manual page x +

man7.org/linux/man 90% Search

Caveats

Do not use **system()** from a privileged program (a set-user-ID or set-group-ID program, or a program with capabilities) because strange values for some environment variables might be used to subvert system integrity. For example, **PATH** could be manipulated so that an arbitrary program is executed with privilege. Use the **exec(3)** family of functions instead, but not **execlp(3)** or **execvp(3)** (which also use the **PATH** environment variable to search for an executable).

system() will not, in fact, work properly from programs with set-user-ID or set-group-ID privileges on systems on which **/bin/sh** is bash version 2: as a security measure, bash 2 drops privileges on startup. (Debian uses a different shell, **dash(1)**, which does not do this when invoked as **sh**.)

```
base distro: Ubuntu 17.10
Kernel:Linux version 4.13.0-32-generic (buildd@lgw01-amd64-016) (gcc version 7.2.0 (Ubuntu
untu3)) #35-Ubuntu SMP Thu Jan 25 09:13:46 UTC 2018
CPU: Intel(R) Core(TM) i7-6500U CPU @ 2.50GHz
Uptime (as of): Mon Feb 19 18:04:12 IST 2018
18:04:12 up 21:31, 1 user, load average: 0.07, 0.03, 0.00
Battery 0: Discharging, 25%, 01:15:45 remaining

Welcome to
Seawolf Mindev

*** IMP NOTE ***
PLEASE READ this file to gain maximum mileage!
~/README_seawolf_mindev_usefulnotes.txt

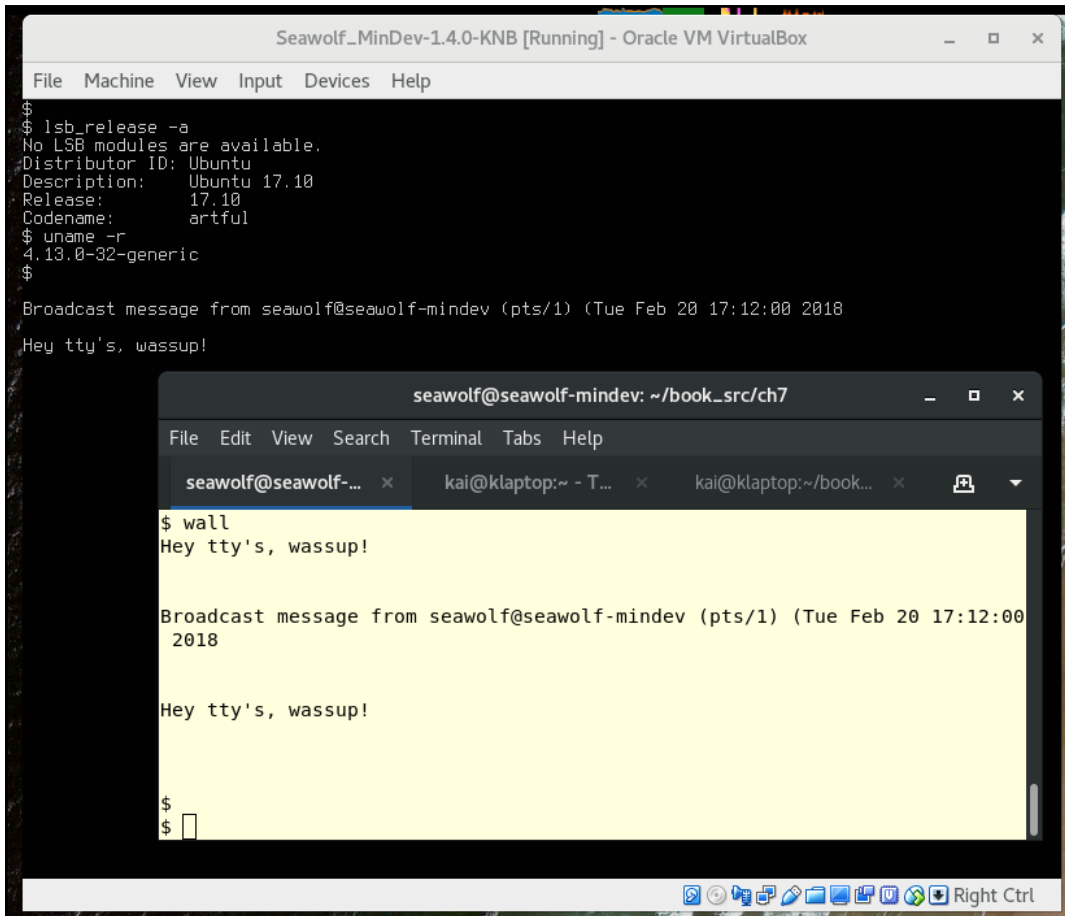
----- To be taken with a pinch of salt :-)- -----
Your reasoning powers are good, and you are a fairly good planner

* Documentation: https://help.ubuntu.com
* Management: https://landscape.canonical.com
* Support: https://ubuntu.com/advantage

Running: ./0setup.bash
IP addr: 192.168.2.20
$ -

[+] Disk Activity LED(10834): initial
[+] Disk Activity LED(10834): init do
```

```
File Edit View Search Terminal Tabs Help
seawolf@seawolf-mindev: ~/book_src/ch7 x
$ ps
PID TTY TIME CMD
6012 pts/0 00:00:00 bash
6714 pts/0 00:00:00 ps
$ id -u
1000
$ ./rootsh_hack2
!WARNING! rootsh_hack2.c:main:36: setuid(0) failed!
perror says: Operation not permitted
seawolf@seawolf-mindev:~/book_src/ch7$ id -u
1000
seawolf@seawolf-mindev:~/book_src/ch7$ ps
PID TTY TIME CMD
6012 pts/0 00:00:00 bash
6724 pts/0 00:00:00 rootsh_hack2
6725 pts/0 00:00:00 sh
6726 pts/0 00:00:00 bash
6750 pts/0 00:00:00 ps
seawolf@seawolf-mindev:~/book_src/ch7$ exit
exit
$
```




```
$ sudo chown root savedset_demo
[sudo] password for seawolf:
$ sudo chmod u+s savedset_demo
$ ls -l savedset_demo
-rwsrwxr-x 1 root seawolf 13144 Feb 20 09:22 savedset_demo*
$ ./savedset_demo
t0: Init:
RUID=1000 EUID=0
RGID=1000 EGID=1000
  Ok, we're effectively running as root! (EUID==0)
t1: Becoming my original self!
RUID=1000 EUID=1000
RGID=1000 EGID=1000
t2: Switching to privileged state now...
RUID=1000 EUID=0
RGID=1000 EGID=1000
  Yup, we're root again!
t3: Switching back to unprivileged state now ...
RUID=1000 EUID=1000
RGID=1000 EGID=1000
$ █
```

Chapter 8: Process Capabilities

Transformation of capabilities during `execve()`

During an `execve(2)`, the kernel calculates the new capabilities of the process using the following algorithm:

```
P'(ambient)      = (file is privileged) ? 0 : P(ambient)

P'(permitted)    = (P(inheritable) & F(inheritable)) |
                  (F(permitted) & cap_bset) | P'(ambient)

P'(effective)    = F(effective) ? P'(permitted) : P'(ambient)


P'(inheritable) = P(inheritable)    [i.e., unchanged]
```

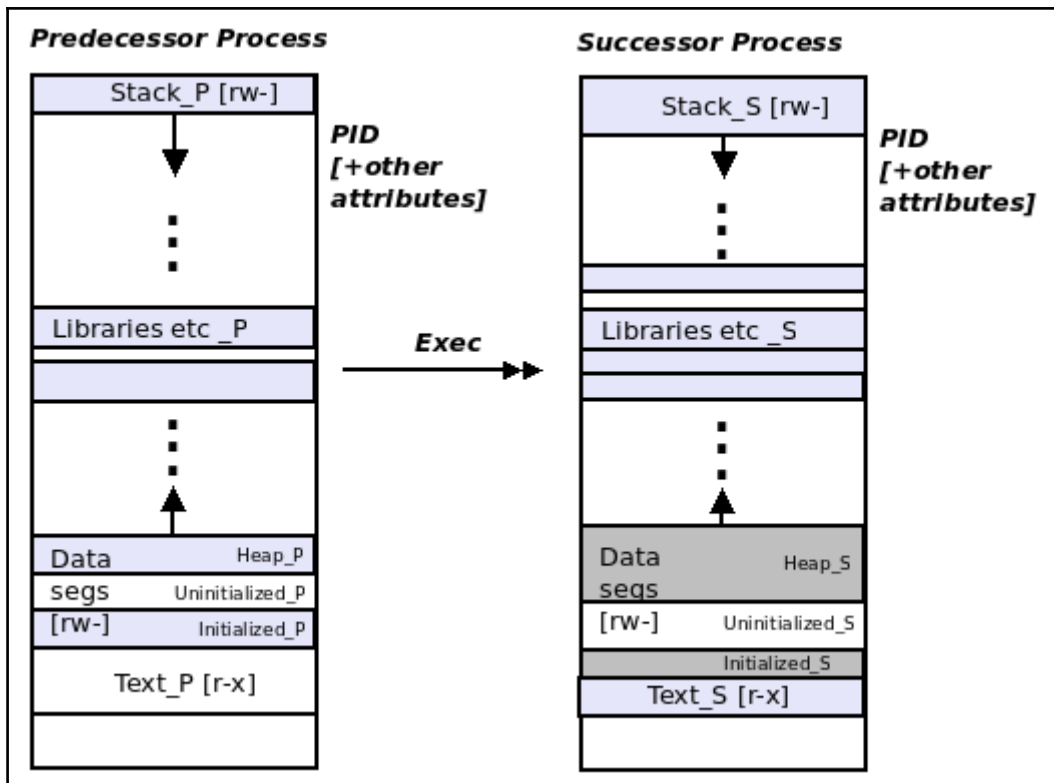
where:

```
P          denotes the value of a thread capability set before the execve(2)
P'         denotes the value of a thread capability set after the execve(2)
F          denotes a file capability set
cap_bset  is the value of the capability bounding set (described below).
```

```
$ ls -l /usr/bin/passwd /usr/bin/write /usr/bin/ping /bin/dumpcap
-rwxr-x---. 1 root wireshark 109344 Jan 19 19:45 /bin/dumpcap
-rwsr-xr-x. 1 root root      27880 Aug  4 2017 /usr/bin/passwd
-rwxr-xr-x. 1 root root      62080 Aug  3 2017 /usr/bin/ping
-rwxr-sr-x. 1 root tty       19584 Sep 22 14:07 /usr/bin/write
$ █
```

Chapter 9: Process Execution

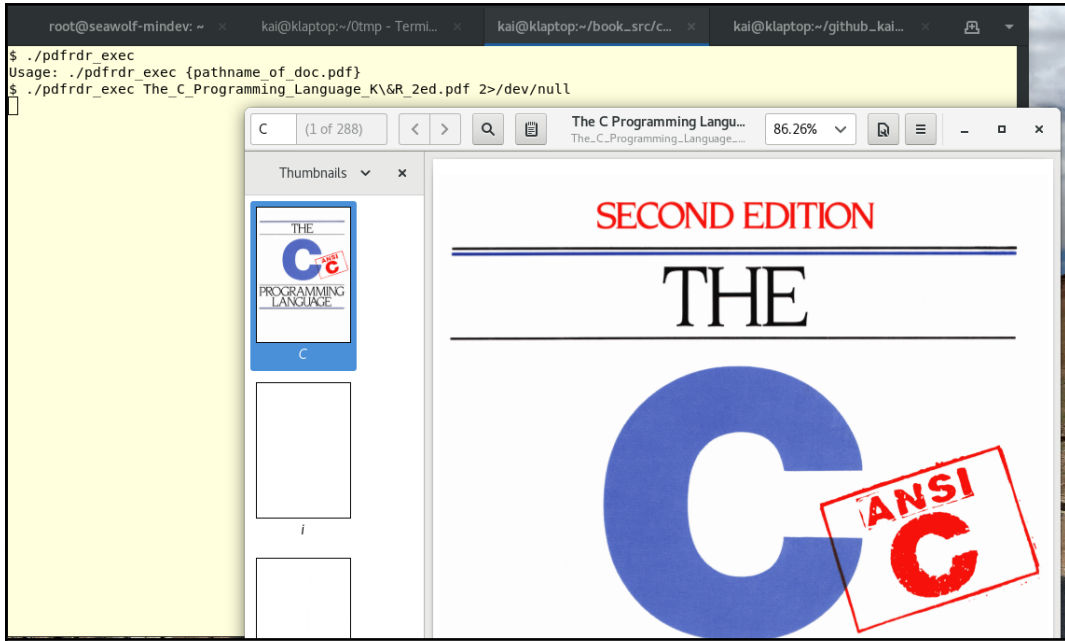




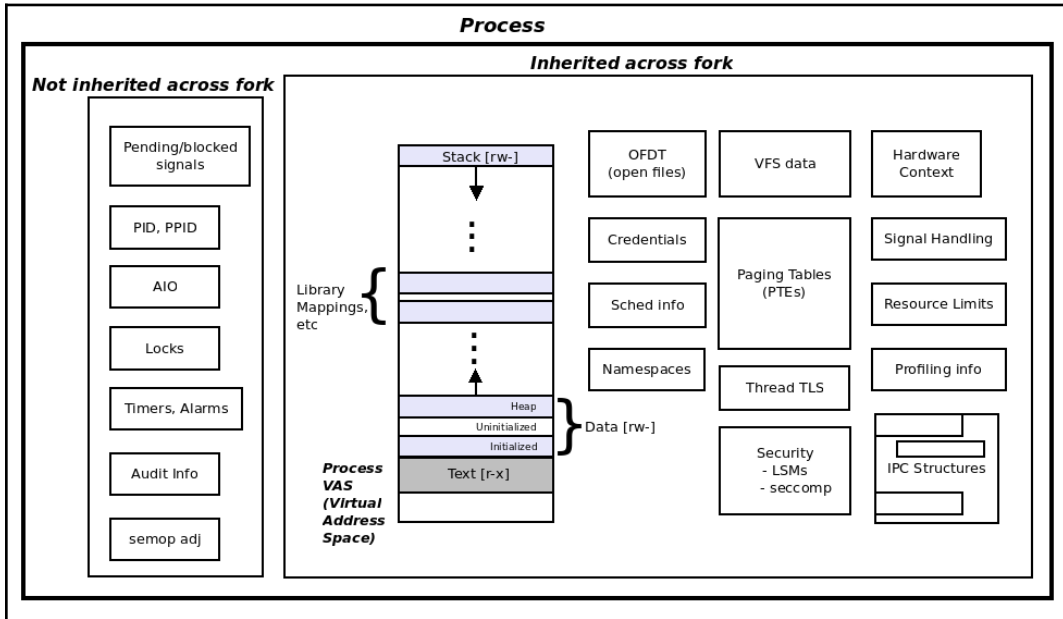
```

$ ps
  PID TTY          TIME CMD
 3396 pts/3        00:00:00 bash
13030 pts/3        00:00:00 ps
$ bash
$ ps
  PID TTY          TIME CMD
 3396 pts/3        00:00:00 bash
13040 pts/3        00:00:00 bash
13087 pts/3        00:00:00 ps
$ exec ps
  PID TTY          TIME CMD
 3396 pts/3        00:00:00 bash
13040 pts/3        00:00:00 ps
$ █

```



Chapter 10: Process Creation





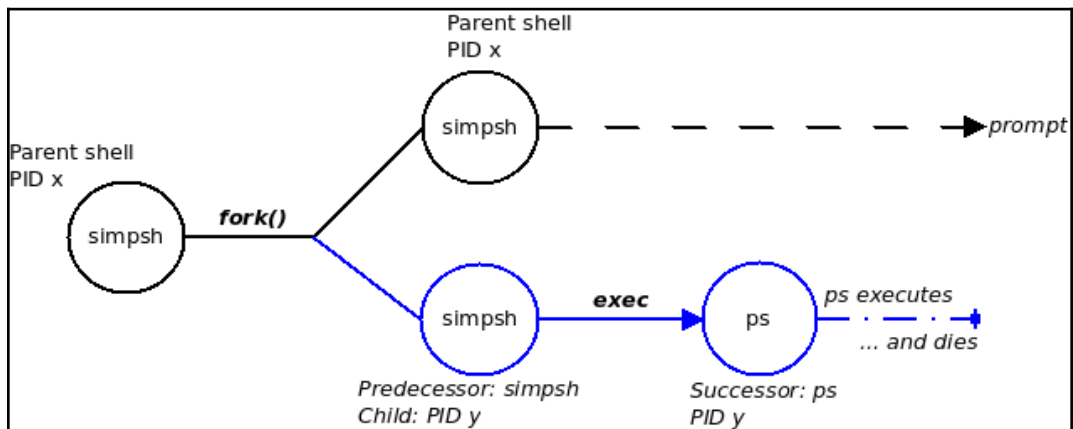

Child Process	Parent Process
<pre> switch((ret = fork())) { case -1 : FATAL("fork failed, aborting!\n"); case 0 : /* Child */ printf("Child process, PID %d:\n" " return %d from fork()\n" , getpid(), ret); foo(atoi(argv[1])); printf("Child process (%d) done.\n", getpid()); exit(EXIT_SUCCESS); default : /* Parent */ printf("Parent process, PID %d:\n" " return %d from fork()\n" , getpid(), ret); bar(atoi(argv[2])); } </pre>	<pre> switch((ret = fork())) { case -1 : FATAL("fork failed, aborting!\n"); case 0 : /* Child */ printf("Child process, PID %d:\n" " return %d from fork()\n" , getpid(), ret); foo(atoi(argv[1])); printf("Child process (%d) done.\n", getpid()); exit(EXIT_SUCCESS); default : /* Parent */ printf("Parent process, PID %d:\n" " return %d from fork()\n" , getpid(), ret); bar(atoi(argv[2])); } </pre>

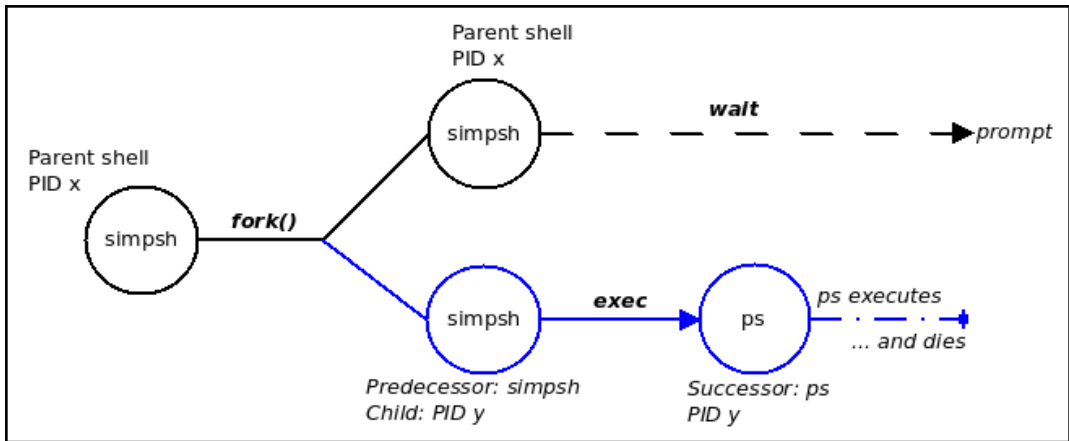
```

switch((ret = fork())) {
  case -1 : FATAL("fork failed, aborting!\n");
  case 0 : /* Child */
    printf("Child process, PID %d:\n"
           " return %d from fork()\n"
           , getpid(), ret);
    foo(atoi(argv[1]));
    printf("Child process (%d) done, exiting ... \n",
           getpid());
    exit(EXIT_SUCCESS);
  default : /* Parent */
    printf("Parent process, PID %d:\n"
           " return %d from fork()\n"
           , getpid(), ret);
    bar(atoi(argv[2]));
}

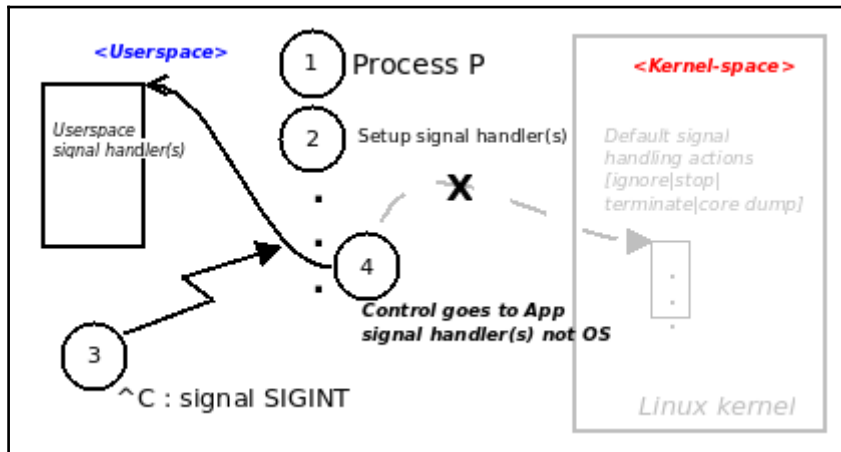
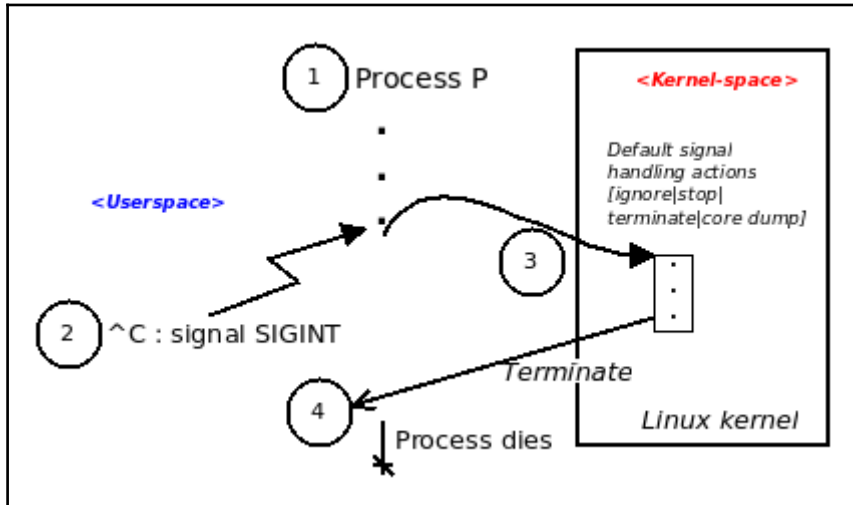
```

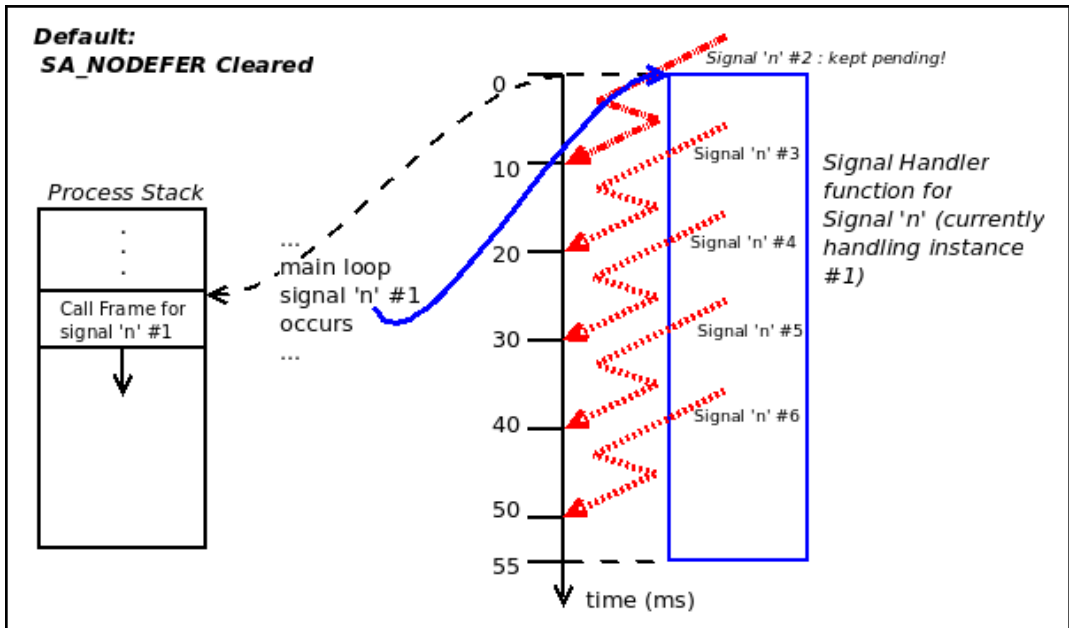








Chapter 11: Signaling - Part I





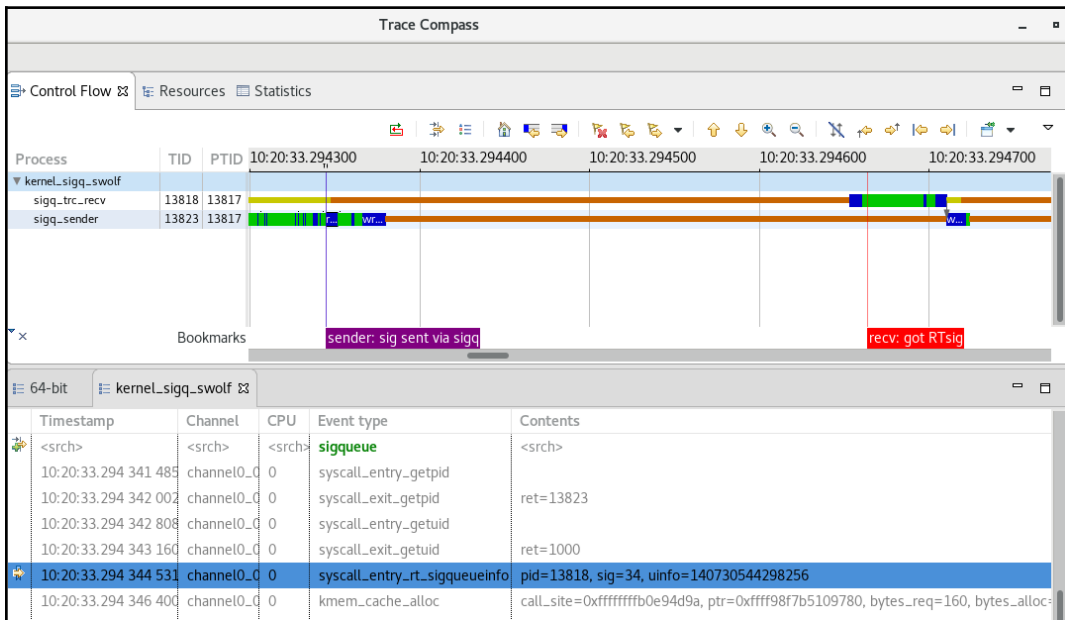


Chapter 12: Signaling - Part II

```

kai@klatp:~$ ./rtsigs_waiter
Trapping the three realtime signals
Process awaiting signals ...
sighdlr: signal 59, s=1; total=1; stack 0x7ffc2f0312c0 : *
sighdlr: signal 39, s=2; total=2; stack 0x7ffc2f0312c0 : *
sighdlr: signal 39, s=3; total=3; stack 0x7ffc2f0312c0 : *
sighdlr: signal 39, s=4; total=4; stack 0x7ffc2f0312c0 : *
sighdlr: signal 39, s=5; total=5; stack 0x7ffc2f0312c0 : *
sighdlr: signal 39, s=6; total=6; stack 0x7ffc2f0312c0 : *
sighdlr: signal 39, s=7; total=7; stack 0x7ffc2f0312c0 : *
sighdlr: signal 39, s=8; total=8; stack 0x7ffc2f0312c0 : *
sighdlr: signal 39, s=9; total=9; stack 0x7ffc2f0312c0 : *
sighdlr: signal 39, s=1; total=10; stack 0x7ffc2f0312c0 : *
sighdlr: signal 39, s=2; total=11; stack 0x7ffc2f0312c0 : *
sighdlr: signal 59, s=3; total=12; stack 0x7ffc2f0312c0 : *
sighdlr: signal 59, s=4; total=13; stack 0x7ffc2f0312c0 : *
sighdlr: signal 59, s=5; total=14; stack 0x7ffc2f0312c0 : *
sighdlr: signal 59, s=6; total=15; stack 0x7ffc2f0312c0 : *
sighdlr: signal 59, s=7; total=16; stack 0x7ffc2f0312c0 : *
sighdlr: signal 59, s=8; total=17; stack 0x7ffc2f0312c0 : *
sighdlr: signal 59, s=9; total=18; stack 0x7ffc2f0312c0 : *
sighdlr: signal 59, s=1; total=19; stack 0x7ffc2f0312c0 : *
sighdlr: signal 59, s=2; total=20; stack 0x7ffc2f0312c0 : *
sighdlr: signal 64, s=3; total=21; stack 0x7ffc2f0312c0 : *
sighdlr: signal 64, s=4; total=22; stack 0x7ffc2f0312c0 : *
sighdlr: signal 64, s=5; total=23; stack 0x7ffc2f0312c0 : *
sighdlr: signal 64, s=6; total=24; stack 0x7ffc2f0312c0 : *
sighdlr: signal 64, s=7; total=25; stack 0x7ffc2f0312c0 : *
sighdlr: signal 64, s=8; total=26; stack 0x7ffc2f0312c0 : *
sighdlr: signal 64, s=9; total=27; stack 0x7ffc2f0312c0 : *
sighdlr: signal 64, s=1; total=28; stack 0x7ffc2f0312c0 : *
sighdlr: signal 64, s=2; total=29; stack 0x7ffc2f0312c0 : *
sighdlr: signal 64, s=3; total=30; stack 0x7ffc2f0312c0 : *
kai@klatp:~$
kai@klatp:~$ ./bombard_sigrt.sh
Usage: bombard_sigrt.sh PID-of-process num-RT-signals-batches-to-send
(-1 to continuously bombard the process with signals).
kai@klatp:~$ ./bombard_sigrt.sh $(pgrep rtsigs_waiter) 10
Sending 10 instances each of RT signal batch
{SIGRTMAX-5, SIGRTMAX, SIGRTMIN+5} to process 3746 ...
i.e. signal #s {59, 64, 39}
SIGRTMAX-5 SIGRTMAX SIGRTMIN+5 SIGRTMAX-5 SIGRTMAX SIGRTMIN+5 SIGRTMAX-5 SIGRTMA
X SIGRTMIN+5 SIGRTMAX-5 SIGRTMAX SIGRTMIN+5 SIGRTMAX-5 SIGRTMAX SIGRTMIN+5 SIGRT
MAX-5 SIGRTMAX SIGRTMIN+5 SIGRTMAX-5 SIGRTMAX SIGRTMIN+5 SIGRTMAX-5 SIGRTMAX SIG
RTMIN+5 SIGRTMAX-5 SIGRTMAX SIGRTMIN+5 SIGRTMAX-5 SIGRTMAX SIGRTMIN+5
kai@klatp:~$

```



Chapter 13: Timers

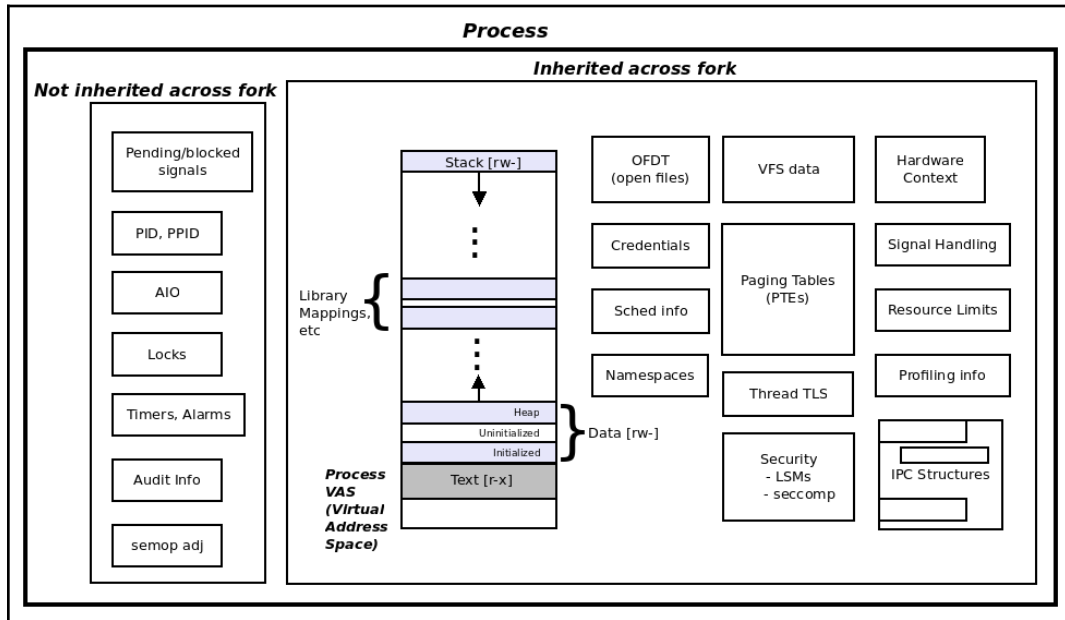
```
$
$
$ ./runwalk_timer 5 2
***** Run Walk Timer *****
          Ver 1.0

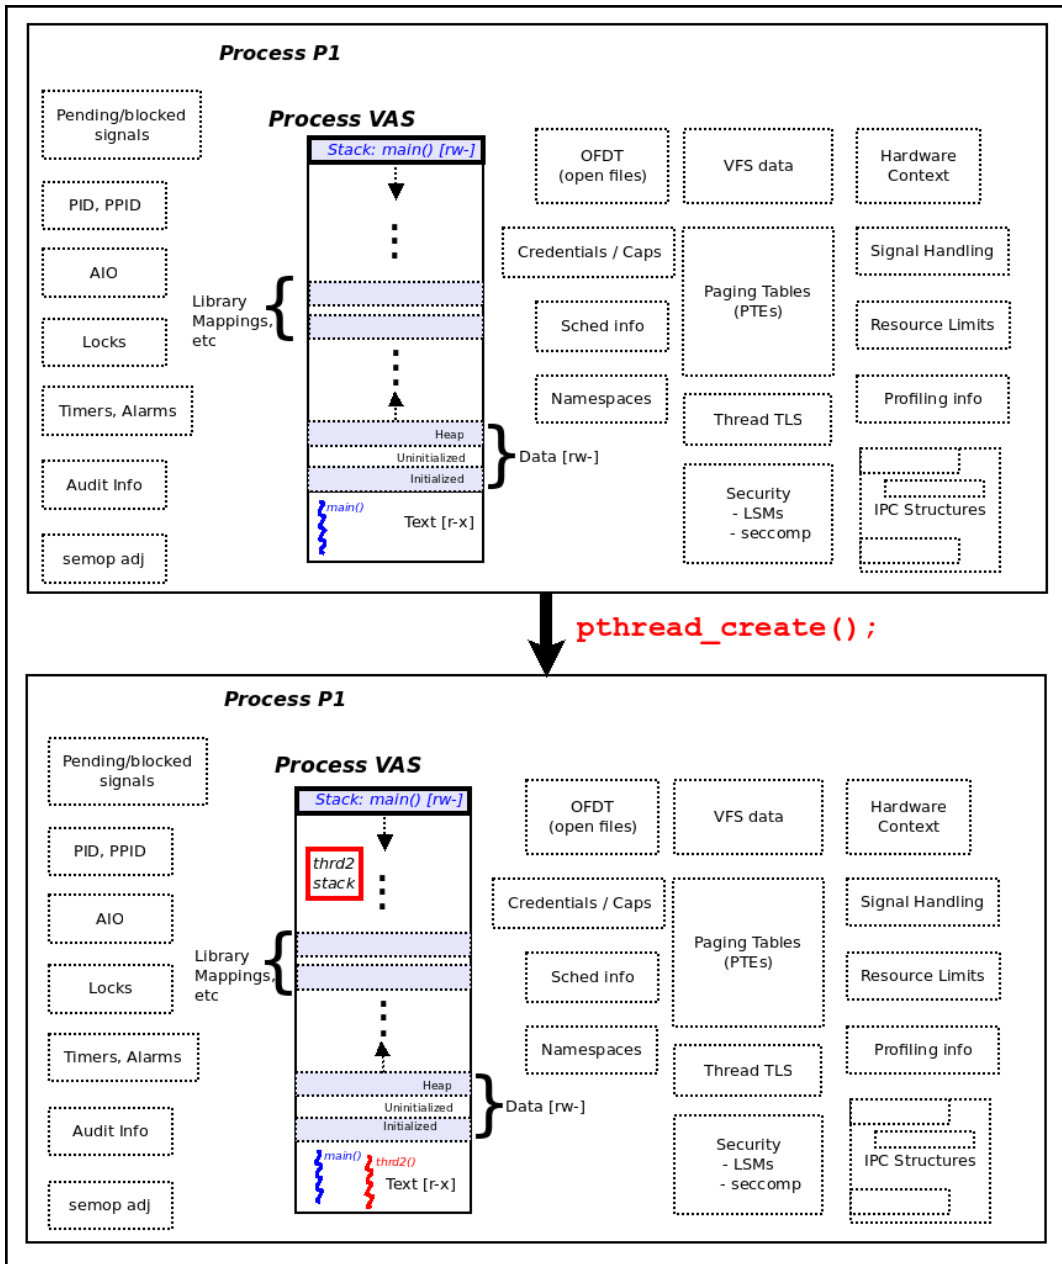
Get moving... Run for 5 seconds
.....
*** Bzzzz!!! WALK! ***      for 2 seconds
..
*** Bzzzz!!! RUN! ***       for 5 seconds
.....
*** Bzzzz!!! WALK! ***      for 2 seconds
..
*** Bzzzz!!! RUN! ***       for 5 seconds
.....
*** Bzzzz!!! WALK! ***      for 2 seconds
..
*** Bzzzz!!! RUN! ***       for 5 seconds
.....
*** Bzzzz!!! WALK! ***      for 2 seconds
..
*** Bzzzz!!! RUN! ***       for 5 seconds

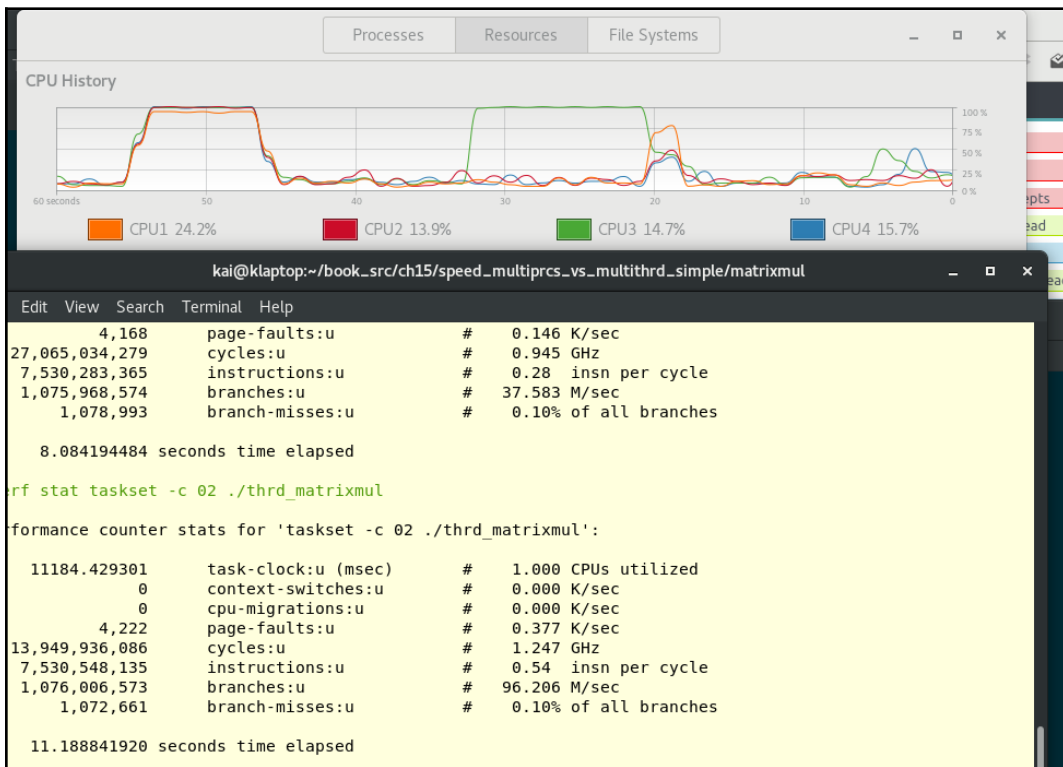
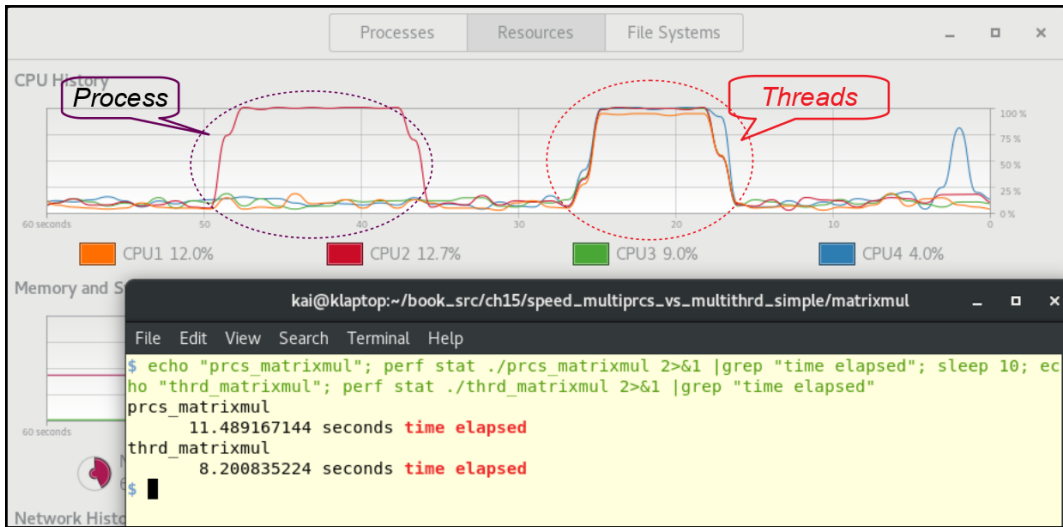
$
$ cat /proc/$(pgrep runwalk_timer)/timers
ID: 0
signal: 34/00007ffd52d4b920
notify: signal/pid.24975
ClockID: 0
$
$
```

Chapter 14: Multithreading with Pthreads

Part I - Essentials









```
seawolf@seawolf-dev-vm: ~
File Edit View Search Terminal Help
$
$
$ ./pthreads3 &
[1] 3906
$ Worker thread #0 running ...
Worker thread #1 running ...
Worker thread #2 running ...

$ ps -LA|grep pthreads3
3906  3906 pts/0    00:00:00 pthreads3 <defunct>
3906  3907 pts/0    00:00:00 pthreads3
3906  3908 pts/0    00:00:00 pthreads3
3906  3909 pts/0    00:00:00 pthreads3
$ #2: work done, exiting now
#1: work done, exiting now
#0: work done, exiting now
```

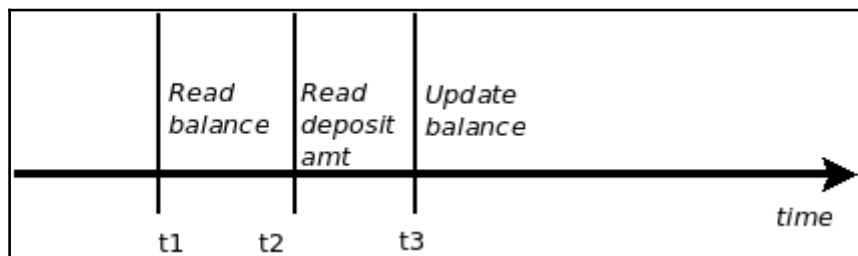


Chapter 15: Multithreading with Pthreads

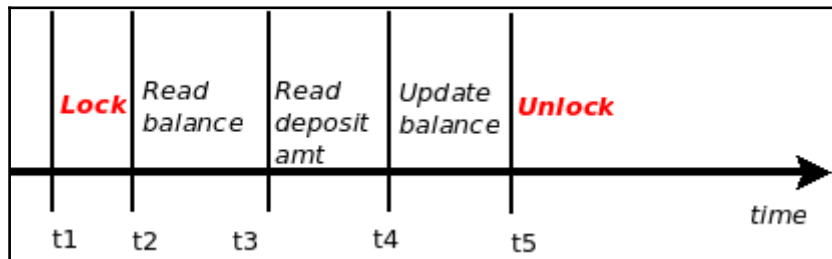
Part II - Synchronization

Time	Thread A on CPU 0	Thread B on CPU 1	Balance
t0			12000
t1	read bal		12000
t2		read bal	12000
t3	deposit 3000		12000
t4		deposit 8000	12000
t5	update bal = 12000+3000		15000
t6		update bal = 15000+8000	23000

Time	Thread A on CPU 0	Thread B on CPU 1	Balance
t0			12000
t1	read bal		12000
t2		read bal	12000
t3	deposit 3000		12000
t4		deposit 8000	12000
t5	update bal = 12000+3000	update bal = 12000+8000	15000
t6			20000



Time	Thread A on CPU 0	Thread B on CPU 1	Balance	
t0			12000	
t1	read bal		12000	Critical Section
t2	deposit 3000		12000	
t3	update bal = 12000+3000		15000	
t4		read bal	15000	Critical Section
t5		deposit 8000	15000	
t6		update bal = 15000+8000	23000	



Secure | <https://godbolt.org>

COMPILER EXPLORER Editor Diff View More

Support diversity in C++ with #include <C++>

C source #1 x x86-64 gcc 8.2 (Editor #1, Compiler #1) C x

```

1 // Type your code here, or load an example.
2 int g=41;
3 int main(void)
4 {
5     g ++;
6 }

```

```

1 g:
2     .long 41
3 main:
4     push rbp
5     mov rbp, rsp
6     mov eax, DWORD PTR g[rip]
7     add eax, 1
8     mov DWORD PTR g[rip], eax
9     mov eax, 0
10    pop rbp
11    ret

```

Output (0/0) gcc (GCC-Explorer-Build) 8.2.0 - 631ms (27768)

Support diversity in C++ with #include

```

1 // Type your code here, or load an example.
2 int g=41;
3 int main(void)
4 {
5     g++;
6 }

```

```

11010 .LX0: .text // ls+ Intel Demangle
1 main:
2     add    DWORD PTR g[rip], 1
3     xor    eax, eax
4     ret
5 g:
6     .long 41

```

Output (0/0) gcc (GCC-Explorer-Build) 8.2.0 - cached (27848)

Support diversity in C++ with #include

```

1 // Type your code here, or load an example.
2 int g=41;
3 int main(void)
4 {
5     g++;
6 }



```


```

11010 .LX0: .text // ls+ Intel Demangle
1 g:
2     .word 41
3 main:
4     str    fp, [sp, #-4]!
5     add    fp, sp, #0
6     ldr    r3, .L3
7     ldr    r3, [r3]
8     add    r3, r3, #1
9     ldr    r2, .L3
10    str    r3, [r2]
11    mov    r3, #0
12    mov    r0, r3
13    add    sp, fp, #0
14    ldr    fp, [sp], #4
15    bx    lr
16 .L3:
17    .word g

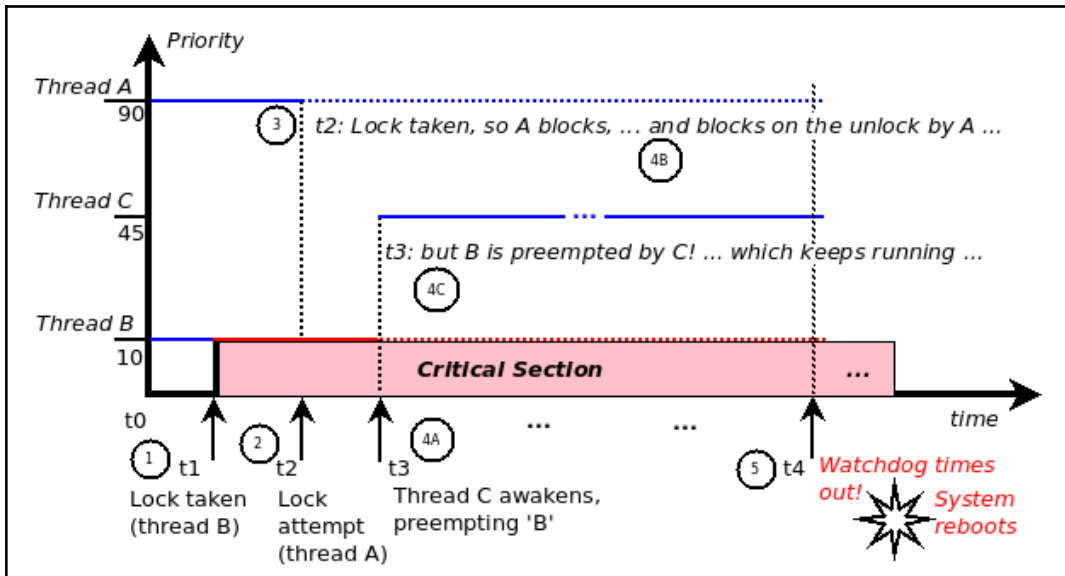
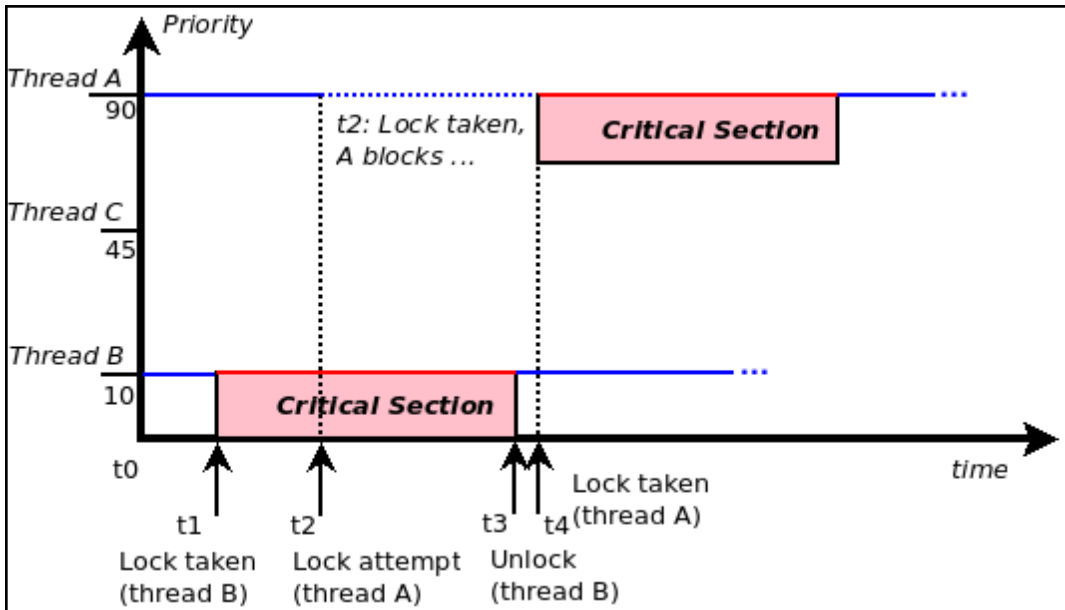
```


Output (0/0) arm-none-eabi-g++ (GNU Tools for Arm Embedded Processors 7-2017-q4-major) 7.2.1 20170904 (release) [ARM/embedded-7-branch revision 255204] - cached





Mutex Type	Robustness	Relock	Unlock When Not Owner
NORMAL	non-robust	deadlock	undefined behavior
NORMAL	robust	deadlock	error returned
ERRORCHECK	either	error returned	error returned
RECURSIVE	either	recursive (see below)	error returned
DEFAULT	non-robust	undefined behavior†	undefined behavior†
DEFAULT	robust	undefined behavior†	error returned





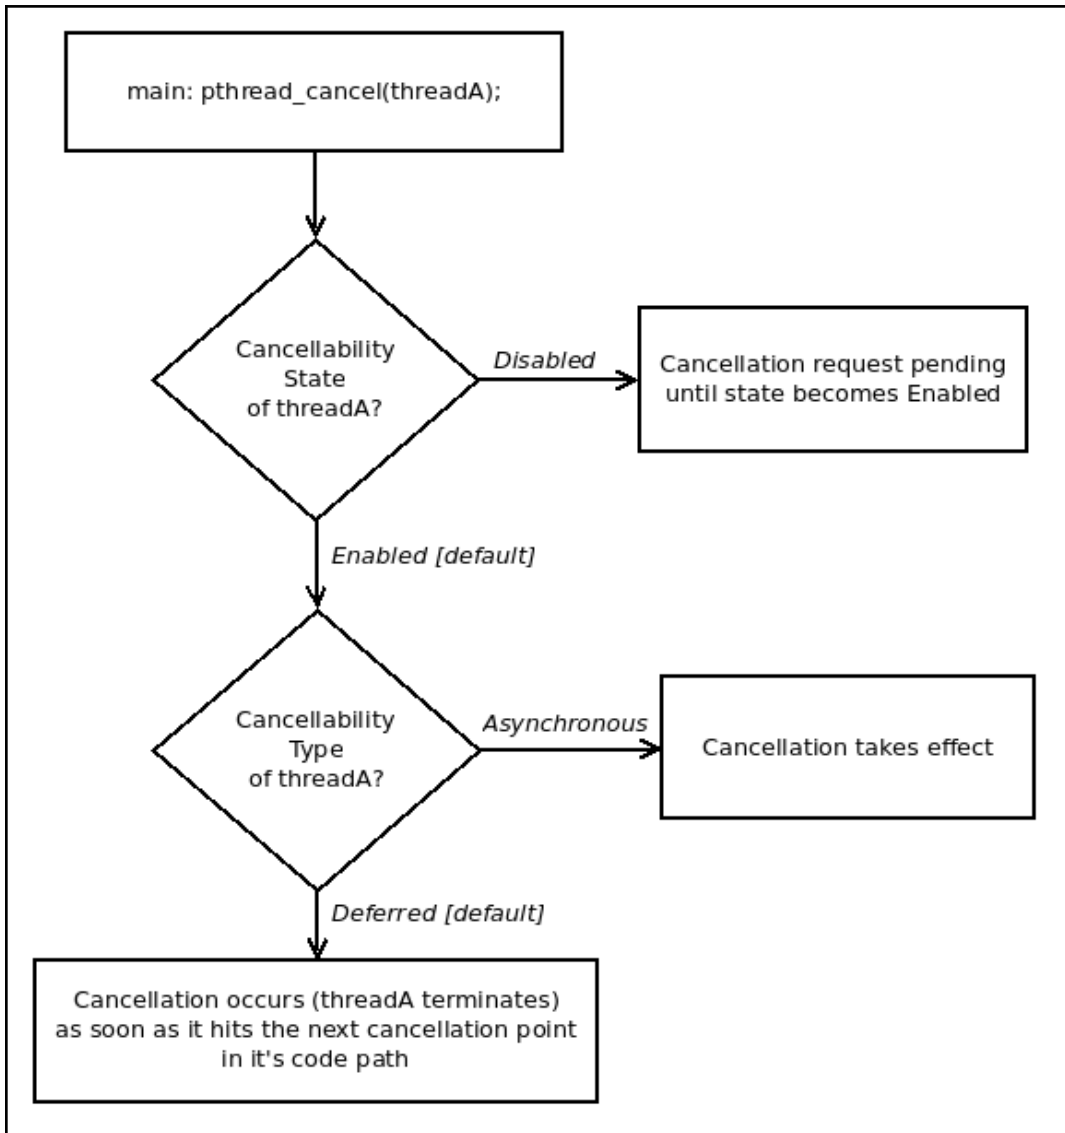
Chapter 16: Multithreading with Pthreads

Part III

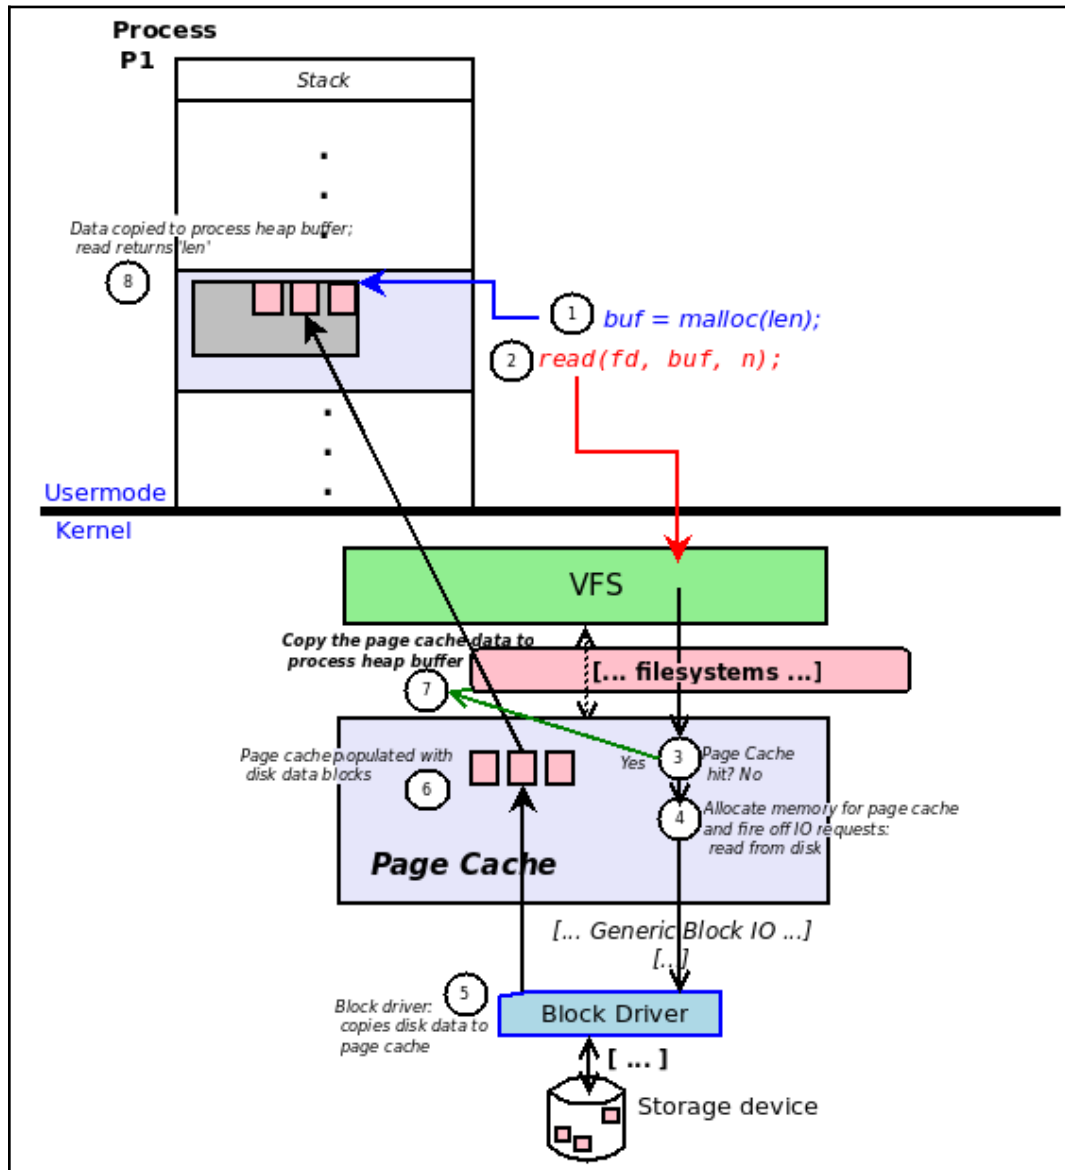
ATTRIBUTES

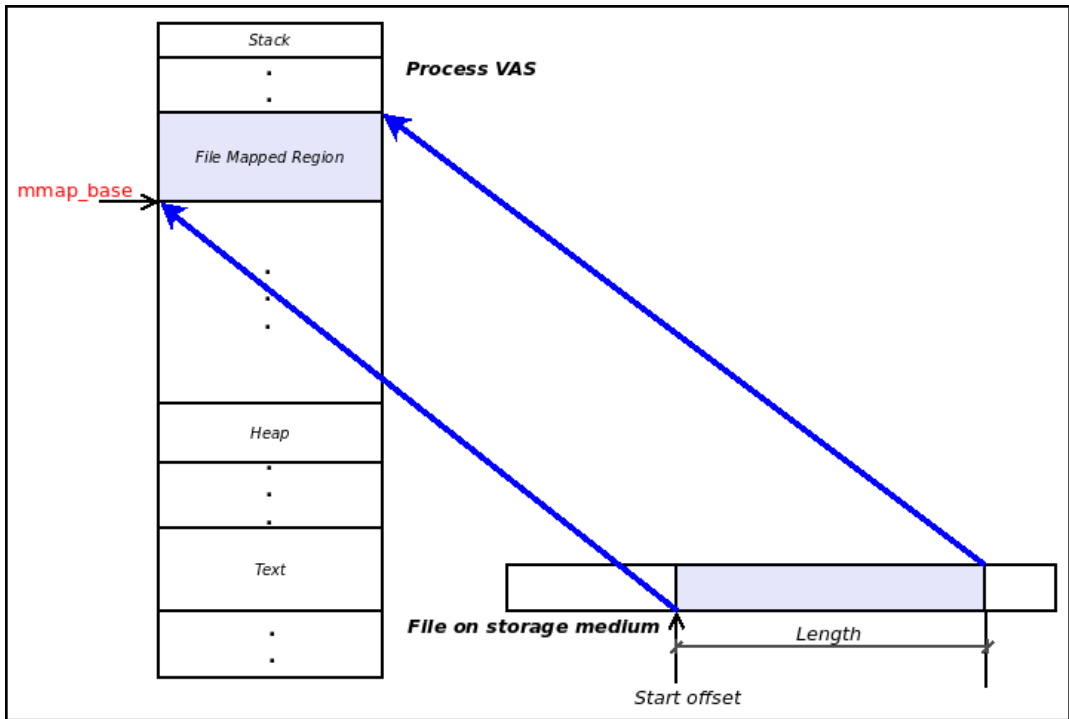
For an explanation of the terms used in this section, see `attributes(7)`.

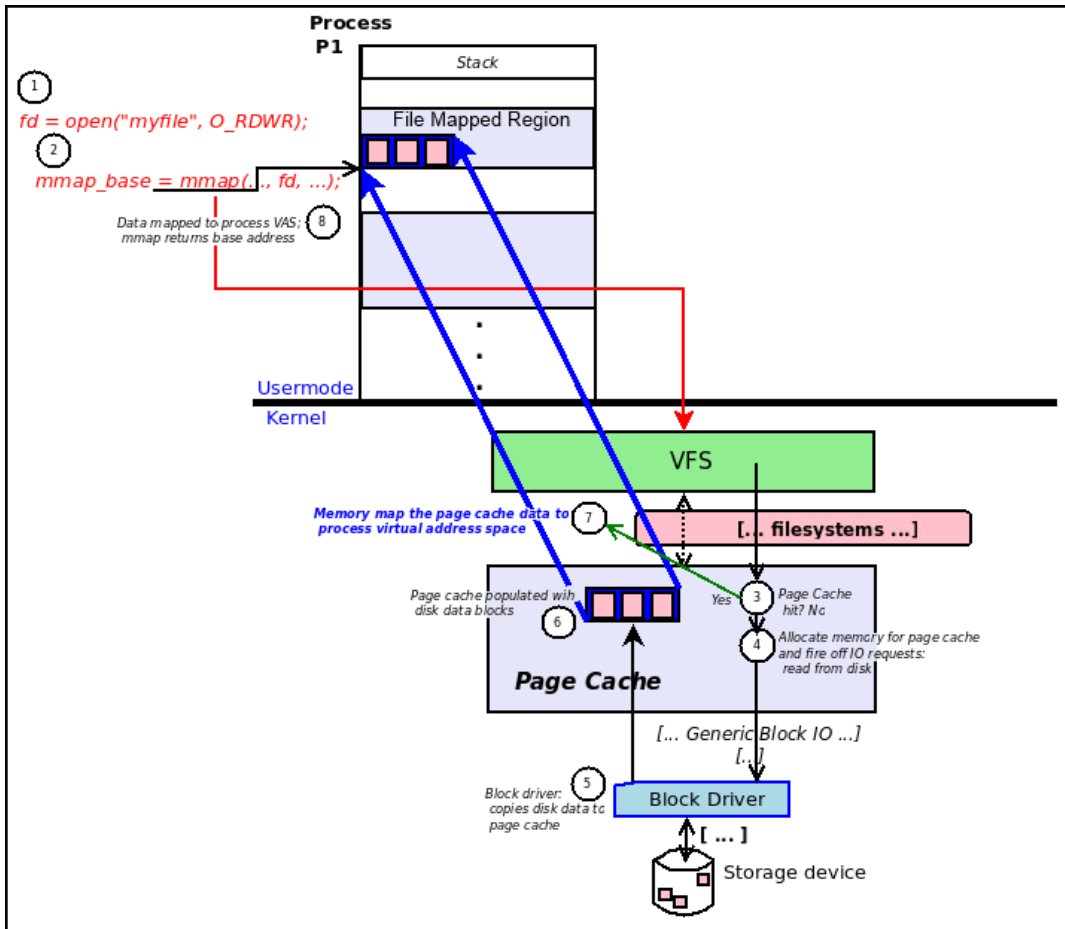
Interface	Attribute	Value
<code>asctime()</code>	Thread safety	MT-Unsafe race:asctime locale
<code>asctime_r()</code>	Thread safety	MT-Safe locale
<code>ctime()</code>	Thread safety	MT-Unsafe race:tmbuf race:asctime env locale
<code>ctime_r()</code> , <code>gmtime_r()</code> , <code>localtime_r()</code> , <code>mktime()</code>	Thread safety	MT-Safe env locale
<code>gmtime()</code> , <code>localtime()</code>	Thread safety	MT-Unsafe race:tmbuf env locale



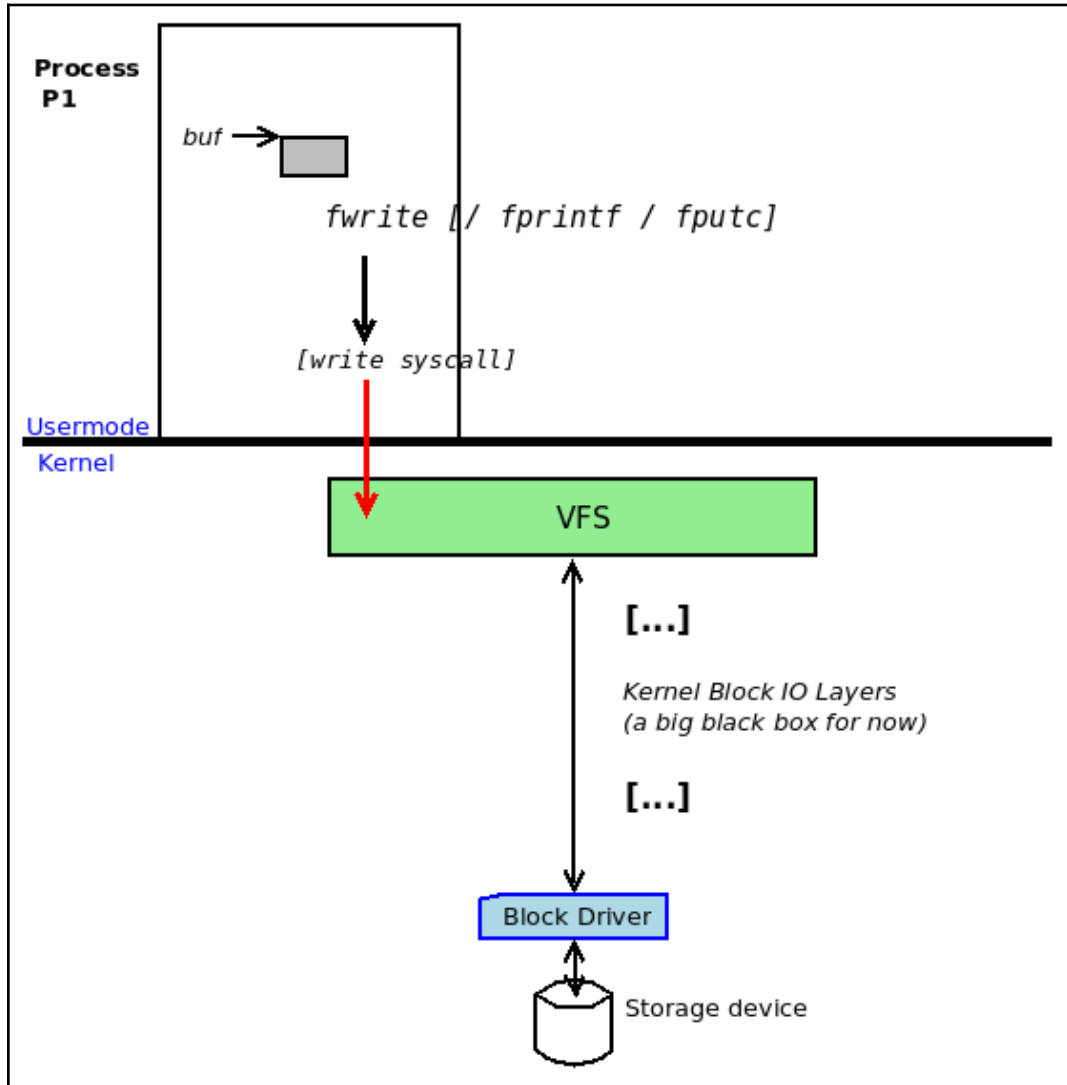
Chapter 18: Advanced File I/O

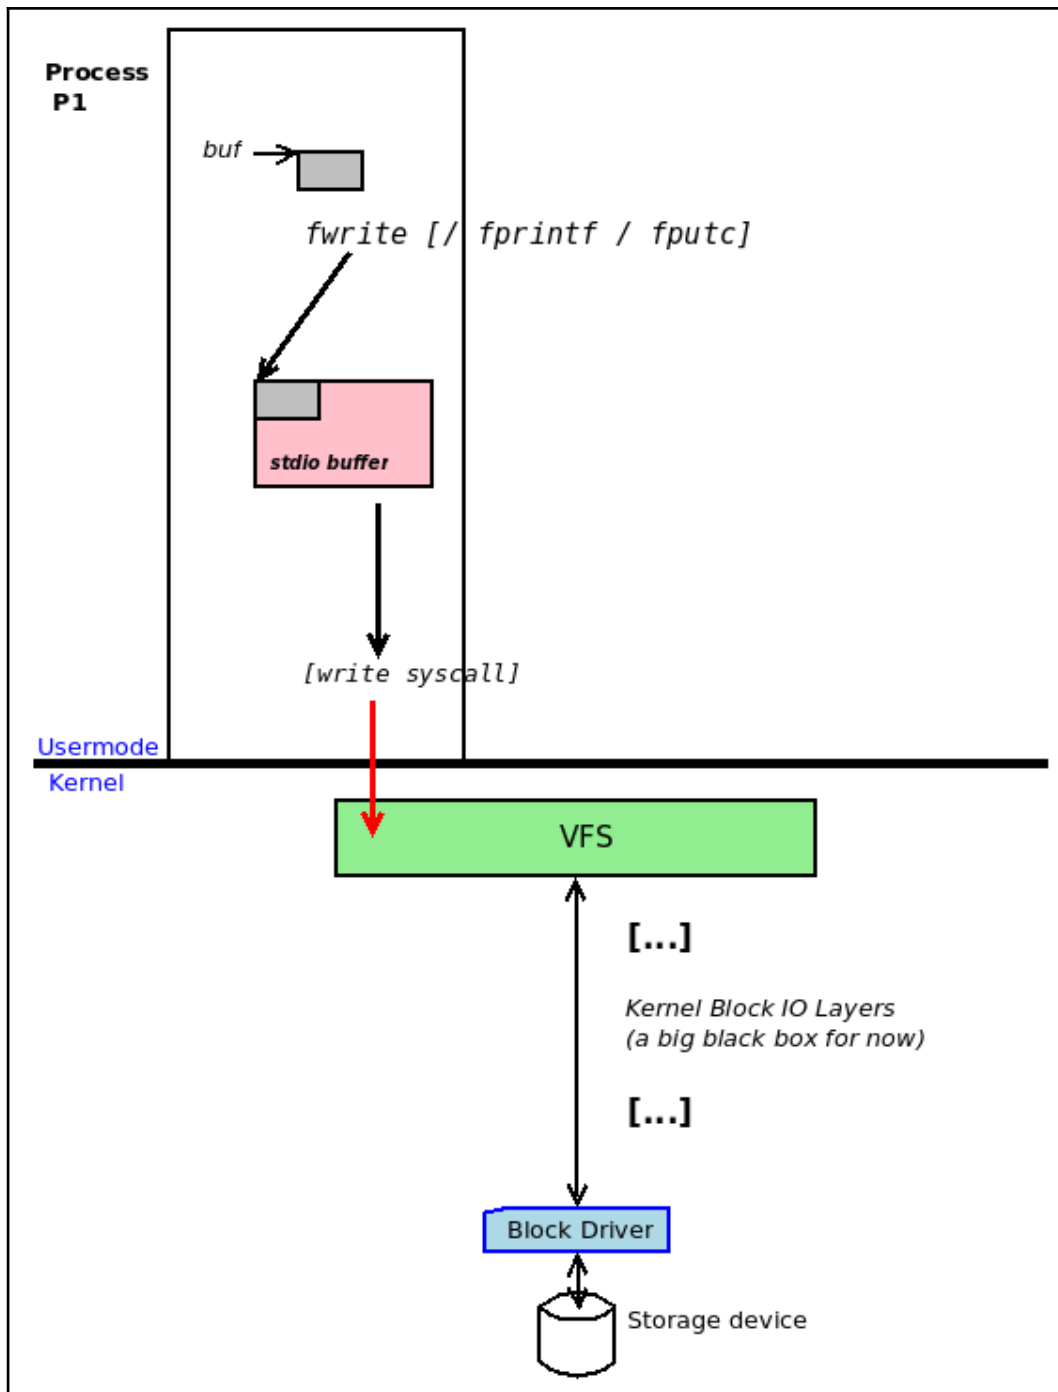




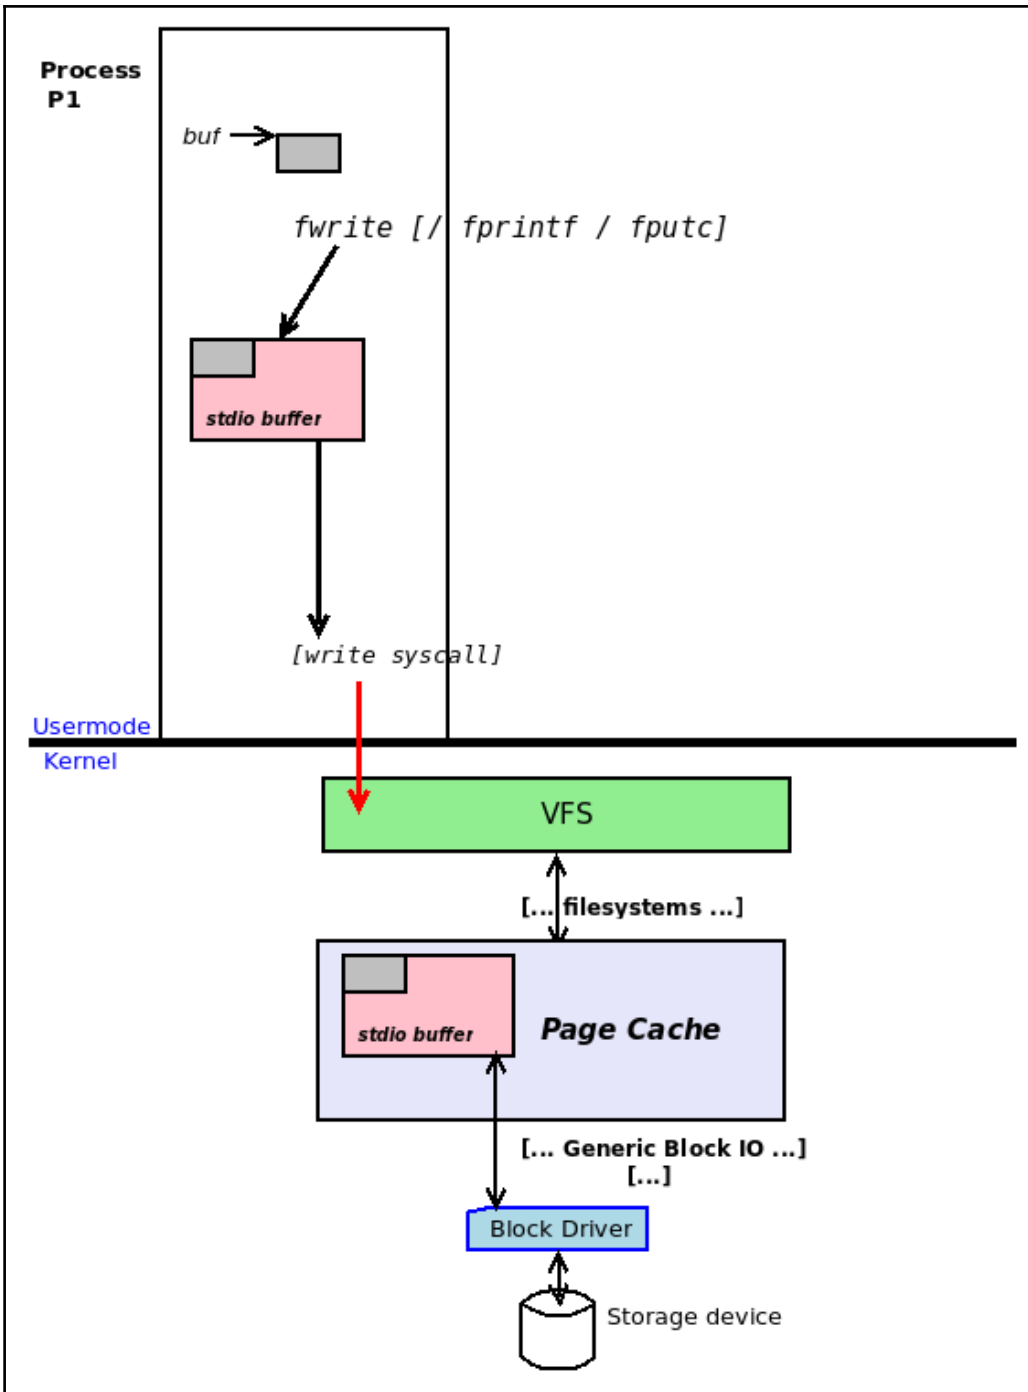


File I/O Essentials

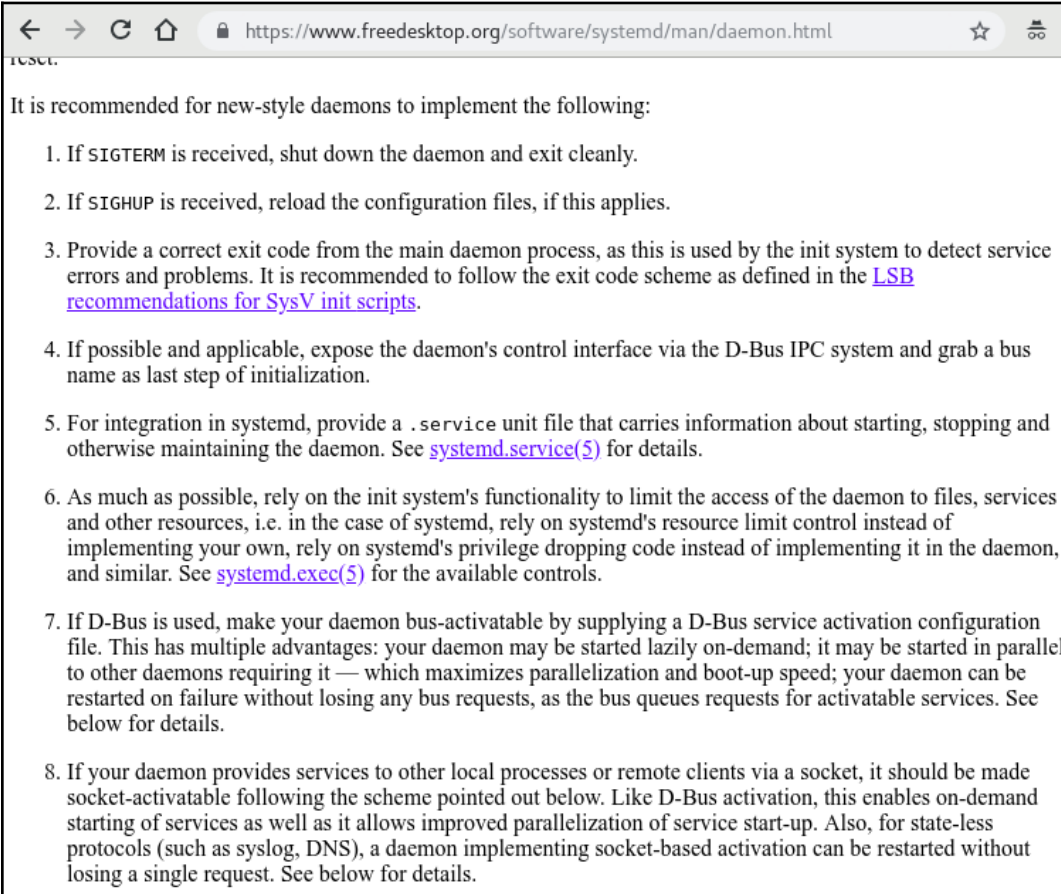




```
$
$ make mode_def
gcc -O2 -Wall -UDEBUG -c mode_def.c -o mode_def.o
mode_def.c: In function 'main':
mode_def.c:31:7: warning: implicit declaration of function 'open'; did you mean
'popen'? [-Wimplicit-function-declaration]
  fd = open(argv[1], O_CREAT|O_RDONLY);
      ^~~~~
      popen
mode_def.c:31:21: error: 'O_CREAT' undeclared (first use in this function)
  fd = open(argv[1], O_CREAT|O_RDONLY);
                    ^~~~~~
mode_def.c:31:21: note: each undeclared identifier is reported only once for each
function it appears in
mode_def.c:31:29: error: 'O_RDONLY' undeclared (first use in this function)
  fd = open(argv[1], O_CREAT|O_RDONLY);
                          ^~~~~~
Makefile:113: recipe for target 'mode_def.o' failed
make: *** [mode_def.o] Error 1
$
```

Daemon Processes



The image is a screenshot of a web browser window. The address bar shows the URL <https://www.freedesktop.org/software/systemd/man/daemon.html>. The page content is titled "daemon(8)" and contains a list of recommendations for new-style daemons. The text is as follows:

It is recommended for new-style daemons to implement the following:

1. If SIGTERM is received, shut down the daemon and exit cleanly.
2. If SIGHUP is received, reload the configuration files, if this applies.
3. Provide a correct exit code from the main daemon process, as this is used by the init system to detect service errors and problems. It is recommended to follow the exit code scheme as defined in the [LSB recommendations for SysV init scripts](#).
4. If possible and applicable, expose the daemon's control interface via the D-Bus IPC system and grab a bus name as last step of initialization.
5. For integration in systemd, provide a `.service` unit file that carries information about starting, stopping and otherwise maintaining the daemon. See [systemd.service\(5\)](#) for details.
6. As much as possible, rely on the init system's functionality to limit the access of the daemon to files, services and other resources, i.e. in the case of systemd, rely on systemd's resource limit control instead of implementing your own, rely on systemd's privilege dropping code instead of implementing it in the daemon, and similar. See [systemd.exec\(5\)](#) for the available controls.
7. If D-Bus is used, make your daemon bus-activatable by supplying a D-Bus service activation configuration file. This has multiple advantages: your daemon may be started lazily on-demand; it may be started in parallel to other daemons requiring it — which maximizes parallelization and boot-up speed; your daemon can be restarted on failure without losing any bus requests, as the bus queues requests for activatable services. See below for details.
8. If your daemon provides services to other local processes or remote clients via a socket, it should be made socket-activatable following the scheme pointed out below. Like D-Bus activation, this enables on-demand starting of services as well as it allows improved parallelization of service start-up. Also, for state-less protocols (such as syslog, DNS), a daemon implementing socket-based activation can be restarted without losing a single request. See below for details.