

# System Programming in Linux

## A Hands-On Introduction

by Stewart N. Weiss

Errata updated to print 1

Page	Error	Correction	Print corrected
58	Linux kernels from 2.6 and later use <b>the sysenter instruction</b> , and the GNU C Library <i>glibc</i> 2.3.2 and later use <b>sysenter</b> .	Linux kernels from 2.6 and later use <b>both the syscall and sysenter instructions</b> , and the GNU C Library <i>glibc</i> 2.3.2 and later use <b>both syscall and sysenter</b> .	Pending
125	<pre>char ch; /* For option handling */</pre>	<pre>int ch; /* For option handling */</pre>	Pending
156	<b>In Figure 4-1, FS3 is missing an arrow.</b>	<p>The diagram illustrates the mapping between user space and kernel space for file management. On the left, two 'Process file descriptor tables' (Process X and Process Y) are shown, each with columns for 'Flags' and 'File ref.'. On the right, two 'System-wide tables' are shown: the 'open file table (file structure table)' and the 'active inode table (per filesystem)'. The open file table has columns for 'File offset', 'Inode ptr.', and 'etc.'. Arrows indicate the flow of information: from Process X and Y descriptors to the open file table, and from the open file table to the active inode table. Specifically, descriptors 0, 1, 2, and m in both process tables point to entries FS 0, FS 1, FS 2, and FS m in the open file table. These entries then point to entries 0, 1, 2, and m in the active inode table. However, descriptor 3 in both process tables has no arrow pointing to the open file table, and entry FS 3 in the open file table has no arrow pointing to the active inode table, which is the error noted in the text.</p> <p>Figure 4-1: The tables used to manage files opened by processes</p>	Pending
239	<pre>BOOL found = FALSE; /* Flag to indicate match found */ char ch;</pre>	<pre>int found = FALSE; /* Flag to indicate match found */ ch;</pre>	Pending

Page	Error	Correction	Print corrected
282	<p>For example, if <code>fd</code> is a valid file descriptor for the file <code>/var/log/lastlog</code>, then</p> <pre>statx(fd, "", AT_EMPTY_PATH, STATX_ALL, &amp;statxbuf);</pre> <p>refers to the file <code>/var/log/lastlog</code>.</p>	<p>For example, if <code>fd</code> is a valid file descriptor for the file <code>/var/log/lastlog</code>, then</p> <pre>statx(fd, "", AT_EMPTY_PATH, STATX_ALL, &amp;statxbuf);</pre> <p>refers to the file <code>/var/log/lastlog</code>.</p>	Pending
286	<pre>int     i; char    ch;</pre>	<pre>int     i; int     ch;</pre>	Pending
307	<pre>\$ ./spl_statfs /media/guest/USB_stick/ File: "/media/guest/USB_stick/" ID: 840 0 Namelen: 1530 Type: msdos</pre>	<pre>\$ ./spl_statfs /media/guest/USB_stick/ File: "/media/guest/USB_stick/" ID: 840000000000 Namelen: 1530 Type: msdos</pre>	Pending
330	<pre>double f(int (*funcp)(int, int), int*);</pre>	<pre>double f(int (*funcp)(int, int), int);</pre>	Pending
335	In Linux and most Unix systems, the hierarchy has no cycles if it has no symbolic links.	In Linux and most Unix systems, the hierarchy <b>usually</b> has no cycles if it has no symbolic links.	Pending
377	<p>For each name that we want to prepend, we get the string length of the name</p> <pre>len = strlen{dirname};</pre> <p>and subtract it from <code>front</code> so that <code>front</code> points to where this name will start in the buffer.</p>	<p>For each name that we want to prepend, we get the string length of the name</p> <pre>len = strlen(dirname);</pre> <p>and subtract it from <code>front</code> so that <code>front</code> points to where this name will start in the buffer.</p>	Pending
378	<pre>while ( (direntp = readdir dir_ptr) != NULL ) {     errno = 0;     get_dev_ino(direntp-&gt;d_name, &amp;this_entry); }</pre>	<pre>while ( (direntp = readdir(dir_ptr) ) != NULL ) {     errno = 0;     get_dev_ino(direntp-&gt;d_name, &amp;this_entry); }</pre>	Pending
391	<b>SIGTRAP</b> Used to implement debuggers and tracing programs such as <code>strace</code> and <code>ptrace</code> .	<b>SIGTRAP</b> Used to implement debuggers <b>such as <code>gdb</code></b> and tracing programs such as <code>strace</code> and <code>ptrace</code> .	Pending
404	<pre>printf("\ncleaning up in progress...\ndone\n"); raise(SIGTSTP); printf("raise() called to stop process.\n");</pre>	<pre>printf("\ncleaning up in progress...\ndone\n"); /* UNSAFE */ raise(SIGTSTP); printf("raise() called to stop process.\n"); /* UNSAFE */</pre>	Pending

Page	Error	Correction	Print corrected
411	If we enter CTRL-C 6 times while it is blocked, when it becomes unblocked, the handler will run just once, not <b>10</b> times.	If we enter CTRL-C 6 times while it is blocked, when it becomes unblocked, the handler will run just once, not <b>6</b> times.	Pending
422	<pre> #include &lt;fenv.h&gt;  void fpe_handler(int signo, siginfo_t *info, void *context) { /* These calls to printf() are all UNSAFE. */   printf("Signal: %s\n", strsignal(info-&gt;si_signo));   switch ( info-&gt;si_code ) {   case FPE_INTDIV:     printf("Code: FPE_INTDIV (Integer divide by zero)\n"); break;   case FPE_FLTDIV:     printf("Code: FPE_FLTDIV (Floating-point divide by zero)\n"); break;   case FPE_FLTOVF:     printf("Code: FPE_FLTOVF (Floating-point overflow)\n"); break;     --snip--   }    raise(SIGTERM); } </pre>	<pre> #include &lt;fenv.h&gt;  <b>BOOL noreturn = FALSE;</b> void fpe_handler(int signo, siginfo_t *info, void *context) { /* These calls to printf() are all UNSAFE. */   printf("Signal: %s\n", strsignal(info-&gt;si_signo));   switch ( info-&gt;si_code ) {   case FPE_INTDIV:     printf("Code: FPE_INTDIV (Integer divide by zero)\n"); break;   case FPE_FLTDIV:     printf("Code: FPE_FLTDIV (Floating-point divide by zero)\n"); break;   case FPE_FLTOVF:     printf("Code: FPE_FLTOVF (Floating-point overflow)\n"); break;     --snip--   }    <b>if ( !noreturn )</b>     raise(SIGTERM); <b>/* To force program to quit */</b> } </pre>	Pending
442	<p><b>The <code>timespec_diff()</code> function introduced on page 442 and used in several programs will not produce a correct time difference when the first time argument represents an older time than the second one. It is intended to be used only when the first is either equal to or later than the second. This is how it is used in all programs that call it. The book's source code repository has another version, <code>timespec_diff2()</code>, that handles this case correctly. The <code>assert.h</code> header file should be included in <code>time_utils.c</code> where the function is defined.</b></p>	<pre> <b>/** timespec_diff(ts1, ts2, *diff) computes the difference ts1 - ts2,     storing it in *diff <b>provided that ts2 is no later than ts1.</b> */</b> </pre> <pre> void timespec_diff(struct timespec ts1, struct timespec ts2,                   struct timespec *diff) {   long temp;   diff-&gt;tv_sec = ts1.tv_sec - ts2.tv_sec;   temp = ts1.tv_nsec - ts2.tv_nsec;   <b>assert ( diff-&gt;tv_sec &gt;= 0 );</b>   ❶ if ( temp &lt; 0 ) {     /* Because temp &lt; 0, we need to borrow 1 sec from tv_sec and        add it to tv_nsec as 1000000000 nanoseconds. */     diff-&gt;tv_sec--;     diff-&gt;tv_nsec = 1000000000 + temp;   }   else     diff-&gt;tv_nsec = temp; } </pre>	Pending

Page	Error	Correction	Print corrected
461	<p>Author's note: There is an extremely tiny probability of a race condition due to the call to <code>draw_initial_bar()</code> preceding the signal handler setup in Listing 9-6. It should be placed after the two calls to <code>sigaction()</code>.</p>	<pre>int main(int argc, char *argv[]) {     struct sigaction    act;     const struct timespec slight_pause = {2, 0}; /* 2-second interval */     struct timespec     remaining_sleep; <del>draw_initial_bar();</del>     sigemptyset(&amp;(act.sa_mask));     act.sa_flags = 0;     act.sa_handler = sig_handler;     if ( sigaction(SIGINT, &amp;act, NULL) == -1 )         fatal_error(errno, "sigaction");     if ( sigaction(SIGQUIT, &amp;act, NULL) == -1 )         fatal_error(errno, "sigaction"); <del>draw_initial_bar();</del>      act.sa_handler = refresh_progressbar; }</pre>	Pending
472	<p>bigfile <span style="color: red;">0923331</span></p>	<p>bigfile <span style="color: red;">923331</span></p>	Pending
486	<pre>struct sigaction sa; char            c;</pre>	<pre>struct sigaction sa; int             c;</pre>	Pending
509	<pre>for ( i = 0; i &lt; elf_header64-&gt;e_phnum; i++ ) {     print_progheader64(fd, &amp;(phdr64[i])); }</pre>	<pre>for ( i = 0; i &lt; elf_header64-&gt;e_phnum; i++ ) {     print_progheader64(fd, &amp;(phdr64[i])); }</pre>	Pending
514	<pre>#include &lt;stdio.h&gt; <del>#include &lt;stdio.h&gt;</del></pre>	<pre>#include &lt;stdio.h&gt;</pre>	Pending
553	<p>In modern Linux systems, the physical pages of the parent's memory image aren't actually copied until <b>the child process</b> attempts to modify them.</p>	<p>In modern Linux systems, the physical pages of the parent's memory image aren't actually copied until <b>one of the processes</b> attempts to modify them.</p>	Pending

Page	Error	Correction	Print corrected
548	<pre>int main(int argc, char *argv[]) {     pid_t newpid;     int    num_children = 4;</pre>	<pre>int main(int argc, char *argv[]) {     pid_t newpid;     int    num_children = 3;</pre>	Pending
581	<pre>\$ ./waitpid_demo Producer PID = 37837 Producer is about to start producing data. Consumer PID = 37838 Data from producer, converted to uppercase: PPPPP Process 37837 stopped by signal 19 # Issued kill -19 37837 from 2nd terminal PPPP Process 37837 continued # Issued kill -18 37837 from 2nd terminal Producer finished; waiting for consumer to finish. Process 37838 stopped by signal 19 # Issued kill -19 37838 from 2nd terminal Process 37838 continued # Issued kill -18 37838 from 2nd terminal PPPPPPPP Process 37838 terminated by signal 8 # Issued kill -8 37838 from 2nd terminal \$</pre>	<pre>\$ ./waitpid_demo Producer PID = 37837 Producer is about to start producing data. Consumer PID = 37838 Data from producer, converted to uppercase: PPPPP Process 37837 stopped by signal 19 # Issued kill -19 37837 from 2nd terminal PPPP Process 37837 continued # Issued kill -18 37837 from 2nd terminal <b>Producer (PID=37837) finished producing data; exiting</b>  <b>Process 37837 exited with status 63</b>  Producer finished; waiting for consumer to finish. Process 37838 stopped by signal 19 # Issued kill -19 37838 from 2nd terminal Process 37838 continued # Issued kill -18 37838 from 2nd terminal PPPPPPPP Process 37838 terminated by signal 8 # Issued kill -8 37838 from 2nd terminal \$</pre>	Pending
582	<p>An entirely different approach is to utilize the SIGCHLD signal. Normally, when a process terminates by calling an exit function or is killed or stopped by a signal, the kernel generates a SIGCHLD signal and sends it to that process's parent, provided that it's still running. There are exceptions to this rule that I'll explain shortly, but for now let's assume that the SIGCHLD signal is always sent to the parent process in these circumstances. In this alternative method, we put the call to wait() or waitpid() inside the SIGCHLD handler.</p>	<p>An entirely different approach is to utilize the SIGCHLD signal. Normally, when a process terminates by calling an exit function or is killed or stopped by a signal, the kernel generates a SIGCHLD signal and sends it to that process's parent, provided that it's still running. There are exceptions to this rule that I'll explain shortly, but for now let's assume that the SIGCHLD signal is always sent to the parent process in these circumstances. In this alternative method, we put the call to wait() or waitpid() inside the SIGCHLD handler.</p>	Pending
614	<pre>t4: P1 executes mov @counter, register1 {counter == 5 }</pre>	<pre>t4: P1 executes mov @counter, register1 {counter == 6 }</pre>	Pending
623	<pre>data = sharedbuf-&gt;buf[sharedbuf-&gt;rear]; sharedbuf-&gt;rear = (sharedbuf-&gt;rear + 1) % BUF_SIZE;</pre>	<pre>data = sharedbuf-&gt;buf[sharedbuf-&gt;rear]; sharedbuf-&gt;front = (sharedbuf-&gt;front + 1) % BUF_SIZE;</pre>	Pending
635	<p>1. Configure a sigevent structure <b>sigev</b> as follows:</p>	<p>1. Configure a sigevent structure <b>sev</b> as follows:</p>	Pending

Page	Error	Correction	Print corrected
636	<pre>struct sigaction sa;</pre>	<pre>struct sigaction sigact;</pre>	Pending
639	<pre>unsigned short int rows, cols; /* Size of terminal window */</pre>	<pre>unsigned short int rows, cols; /* Size of terminal window */ msgtype          newmsg;      /* Message struct for message */</pre>	Pending
658	If such a function existed, we could copy the pipe's write end descriptor into the <b>descriptors</b> for standard <b>input and</b> output.	If such a function existed, we could copy the pipe's write end descriptor into the <b>descriptor</b> for standard output.	Pending
707	<ol style="list-style-type: none"> <li>2. Modify <i>calc_client.c</i> so that if it has no command line argument, it reads the expression from standard input.</li> <li>3. Modify <i>spl_calc.b</i>, <i>calc_client.c</i>, and <i>calc_server.c</i> so that if the user supplies a <code>-l</code> command line option to the client, the server will request <code>bc</code> to load its standard math library, as described on its man page.</li> <li>4. Modify both <i>calc_client.c</i> and <i>calc_server.c</i> so that if the user supplies an <code>-s</code> value command line option, the number of digits to the right of the decimal point for all answers will be the supplied value.</li> <li>5. Write a version of <i>calc_server.c</i> that still relies upon the <code>bc</code> command but does not call <code>popen()</code> to run it.</li> <li>6. Add a <code>-d</code> option to <i>calc_server.c</i> that, when present, turns it into a daemon, and when not, requires the user to run it in the background.</li> </ol>	<ol style="list-style-type: none"> <li>2. Modify <i>spl_calc_client.c</i> so that if it has no command line argument, it reads the expression from standard input.</li> <li>3. Modify <i>spl_calc.b</i>, <i>spl_calc_client.c</i>, and <i>spl_calc_server.c</i> so that if the user supplies a <code>-l</code> command line option to the client, the server will request <code>bc</code> to load its standard math library, as described on its man page.</li> <li>4. Modify both <i>spl_calc_client.c</i> and <i>spl_calc_server.c</i> so that if the user supplies an <code>-s</code> value command line option, the number of digits to the right of the decimal point for all answers will be the supplied value.</li> <li>5. Write a version of <i>spl_calc_server.c</i> that still relies upon the <code>bc</code> command but does not call <code>popen()</code> to run it.</li> <li>6. Add a <code>-d</code> option to <i>spl_calc_server.c</i> that, when present, turns it into a daemon, and when not, requires the user to run it in the background.</li> </ol>	Pending
722	<pre>thread_data[t].task_id = t;</pre>	<pre>thread_data[t].id = t;</pre>	Pending
785	<pre>BOOL          reader_preference = FALSE; char          ch;</pre>	<pre>BOOL          reader_preference = FALSE; int           ch;</pre>	Pending

Page	Error	Correction	Print corrected																								
814	<pre> while ( more_data_to_read ) {           /* While source file has more data */ // if ( aio_read() completed ) { num_read = aio_return(&amp;aio_block); /* Get num bytes read. */ writebuf = aio_block.aio_buf; /* write() from buffer just filled */ i = i^1; /* Same as i = (i == 0) ? 1 : 0 */ aio_block.aio_buf = buffer[i]; /* Use other buffer for next read. */ aio_block.aio_offset += num_bytes_read; /* Adjust offset for read. */ aio_read(&amp;aio_cb)); /* Request next read. */ write(target_fd, writebuf, num_bytes_read); /* Start write(). */ } } </pre>	<pre> while ( more_data_to_read ) {           /* While source file has more data */ if ( aio_read() completed ) { num_read = aio_return(&amp;aio_block); /* Get num bytes read. */ writebuf = aio_block.aio_buf; /* write() from buffer just filled */ i = i^1; /* Same as i = (i == 0) ? 1 : 0 */ aio_block.aio_buf = buffer[i]; /* Use other buffer for next read. */ aio_block.aio_offset += num_bytes_read; /* Adjust offset for read.*/ aio_read(&amp;aio_cb)); /* Request next read. */ write(target_fd, writebuf, num_bytes_read); /* Start write(). */ } } </pre>	Pending																								
825	<pre> for ( int i = 0; i &lt; NSENDERS; i++ ) { if ( pipe(pipefd[i]) == -1 ) /* Create pipe[pipefd[i]. */ </pre>	<pre> for ( int i = 0; i &lt; NSENDERS; i++ ) { if ( pipe(pipefd[i]) == -1 ) /* Create pipe[pipefd[i]]. */ </pre>	Pending																								
837	<table border="0" style="width: 100%;"> <tr> <td style="text-align: center; width: 50%;">Terminal line buffer</td> <td style="text-align: center; width: 50%;">read() buffer</td> </tr> <tr> <td style="border: 1px solid black; padding: 2px;">abcde</td> <td style="background-color: #cccccc; text-align: center; padding: 2px;">&lt;empty&gt;</td> </tr> <tr> <td style="border: 1px solid black; padding: 2px;">abcde&lt;BS&gt;&lt;BS&gt;</td> <td style="background-color: #cccccc; text-align: center; padding: 2px;">&lt;empty&gt;</td> </tr> <tr> <td style="border: 1px solid black; padding: 2px;">abcde&lt;BS&gt;&lt;BS&gt;xyz</td> <td style="background-color: #cccccc; text-align: center; padding: 2px;">&lt;empty&gt;</td> </tr> <tr> <td style="border: 1px solid black; padding: 2px;">abcde&lt;BS&gt;&lt;BS&gt;xyz&lt;CR&gt;</td> <td style="background-color: #cccccc; text-align: center; padding: 2px;">&lt;empty&gt;</td> </tr> <tr> <td style="border: 1px solid black; padding: 2px; text-align: center;">&lt;empty&gt;</td> <td style="background-color: #cccccc; text-align: center; padding: 2px;">abcxyz</td> </tr> </table> <p>Figure 18-1: A schematic representation of the processing of an edited input line</p>	Terminal line buffer	read() buffer	abcde	<empty>	abcde<BS><BS>	<empty>	abcde<BS><BS>xyz	<empty>	abcde<BS><BS>xyz<CR>	<empty>	<empty>	abcxyz	<table border="0" style="width: 100%;"> <tr> <td style="text-align: center; width: 50%;">Terminal line buffer</td> <td style="text-align: center; width: 50%;">read() buffer</td> </tr> <tr> <td style="border: 1px solid black; padding: 2px;">abcde</td> <td style="background-color: #cccccc; text-align: center; padding: 2px;">&lt;empty&gt;</td> </tr> <tr> <td style="border: 1px solid black; padding: 2px;">abcde&lt;BS&gt;&lt;BS&gt;</td> <td style="background-color: #cccccc; text-align: center; padding: 2px;">&lt;empty&gt;</td> </tr> <tr> <td style="border: 1px solid black; padding: 2px;">abcde&lt;BS&gt;&lt;BS&gt;xyz</td> <td style="background-color: #cccccc; text-align: center; padding: 2px;">&lt;empty&gt;</td> </tr> <tr> <td style="border: 1px solid black; padding: 2px;">abcde&lt;BS&gt;&lt;BS&gt;xyz&lt;CR&gt;</td> <td style="background-color: #cccccc; text-align: center; padding: 2px;">&lt;empty&gt;</td> </tr> <tr> <td style="border: 1px solid black; padding: 2px; text-align: center;">&lt;empty&gt;</td> <td style="background-color: #cccccc; text-align: center; padding: 2px;">abcxyz&lt;CR&gt;</td> </tr> </table> <p>Figure 18-1: A schematic representation of the processing of an edited input line</p>	Terminal line buffer	read() buffer	abcde	<empty>	abcde<BS><BS>	<empty>	abcde<BS><BS>xyz	<empty>	abcde<BS><BS>xyz<CR>	<empty>	<empty>	abcxyz<CR>	Pending
Terminal line buffer	read() buffer																										
abcde	<empty>																										
abcde<BS><BS>	<empty>																										
abcde<BS><BS>xyz	<empty>																										
abcde<BS><BS>xyz<CR>	<empty>																										
<empty>	abcxyz																										
Terminal line buffer	read() buffer																										
abcde	<empty>																										
abcde<BS><BS>	<empty>																										
abcde<BS><BS>xyz	<empty>																										
abcde<BS><BS>xyz<CR>	<empty>																										
<empty>	abcxyz<CR>																										
850	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 20%;">Mask name</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>NOFLSH</td> <td>Disable flushing the input and output queues when generating signals for the <b>INTR</b>, <b>QUIT</b>, and <b>SUSP</b> characters.</td> </tr> </tbody> </table>	Mask name	Description	NOFLSH	Disable flushing the input and output queues when generating signals for the <b>INTR</b> , <b>QUIT</b> , and <b>SUSP</b> characters.	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 20%;">Mask name</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>NOFLSH</td> <td>Disable flushing the input and output queues when generating signals for the <b>INTR</b>, <b>QUIT</b>, and <b>SUSP</b> characters.</td> </tr> </tbody> </table>	Mask name	Description	NOFLSH	Disable flushing the input and output queues when generating signals for the <b>INTR</b> , <b>QUIT</b> , and <b>SUSP</b> characters.	Pending																
Mask name	Description																										
NOFLSH	Disable flushing the input and output queues when generating signals for the <b>INTR</b> , <b>QUIT</b> , and <b>SUSP</b> characters.																										
Mask name	Description																										
NOFLSH	Disable flushing the input and output queues when generating signals for the <b>INTR</b> , <b>QUIT</b> , and <b>SUSP</b> characters.																										

Page	Error	Correction	Print corrected
851	<pre>\$ bash --no-editing</pre>	<pre>\$ bash --noediting</pre>	Pending
855	<p style="text-align: center;"><b>Handling Excess Input in the Terminal</b></p> <p><b>When</b> <code>read()</code> reads from a terminal, it doesn't receive characters until a newline is entered. Assuming that <code>ttyfd</code> refers to a terminal, if a program calls <code>read(ttyfd, &amp;buffer, 32)</code> but the user enters 40 characters and then a newline, <code>read()</code> will receive the first 32 characters and store them into <code>buffer</code>, but there will be 8 characters and a newline remaining in the input queue. A second call to <code>read(ttyfd, &amp;buffer, 32)</code> will be satisfied immediately because <b>there's a newline in the input queue</b>. The remaining 8 characters and the newline will be stored into <code>buffer</code>. If a program doesn't want these characters to be read into <code>buffer</code>, it can call <code>tcflush(ttyfd, TCIFLUSH)</code> to remove them from the input queue before reading again.</p>	<p style="text-align: center;"><b>Handling Excess Input in the Terminal</b></p> <p><b>In canonical mode, when</b> <code>read()</code> reads from a terminal, it doesn't receive characters until a newline is entered <b>(or the user terminates input by entering the end-of-file character)</b>. Assuming that <code>ttyfd</code> refers to a terminal, if a program calls <code>read(ttyfd, &amp;buffer, 32)</code> but the user enters 40 characters and then a newline, <code>read()</code> will receive the first 32 characters and store them into <code>buffer</code>, but there will be 8 characters and a newline remaining in the <b>terminal's</b> input queue. A second call to <code>read(ttyfd, &amp;buffer, 32)</code> will be satisfied immediately because <b>the input queue is not empty</b>. The remaining 8 characters and the newline will be stored into <code>buffer</code>. If a program doesn't want these characters to be read into <code>buffer</code>, it can call <code>tcflush(ttyfd, TCIFLUSH)</code> to remove them from the input queue before reading again.</p>	Pending
874	<i>fillqueue.c</i>	<i>fillqueue.c</i>	Pending
875	In the second terminal, turn off canonical mode and run <code>fillqueue</code> :	In the second terminal, turn off canonical mode and run <code>fillqueue</code> <b>(you may have to set <code>cap_sys_admin</code> on <code>fillqueue</code> as you did for <code>fakeinput_demo</code>):</b>	Pending
933	<pre>    }     wrefresh(win); }</pre>	<pre>    }     wrefresh(win); }</pre>	Pending
936	<pre>if ( *proclist != NULL )</pre>	<pre>if ( *proclist != NULL ) {</pre>	Pending

## Additional Notes

### Page 197

An alternative to the command

```
cpp -v /dev/null -o /dev/null
```

is

```
cpp -v /dev/null -P
```

which generates the same output.

---

### Page 411

The code following the `while` loop in *sigprocmask\_demo1.c* should never be reached. It is there to show how to use `SIG_SETMASK` to set the blocked signals to a previous signal set. It should have a preceding comment noting that it's unreachable as long as the `while` loop cannot be exited.

---

### Page 423

Although it is not incorrect, the extra parentheses in this line

```
sigemptyset(&(action.sa_mask));
```

are unnecessary, because the `struct` `.` operator has higher precedence than the address-of operator, `&`. It can also be:

```
sigemptyset(&action.sa_mask);
```

---

### Page 485

The *posix\_timer\_demo1.c* program relies on the assumption that the timer IDs returned by successive calls to `timer_create()` start at 0 and increase by 1. POSIX doesn't require this. In the versions of Linux on which I have tested this, they do behave this way. An actual application should not be based on this assumption.

---

### Page 611

The two lines

```
shmp->ptr1 = (char*)  
malloc(20);
```

are correct, but it is a single assignment that should be:

```
shmp->ptr1 = (char*) malloc(20);
```

---

### Page 612

Although it is not incorrect, the line

```
printf("%s\n", (char*) (shmp->ptr2));
```

should be

```
printf("%s\n", shmp->ptr2);
```

since `shmp->ptr2` is a string pointer (`char*`).

---

### Page 618

In *unnamedsem\_demo.c*, although this assignment is correct:

```
shmp = mmap(NULL,
            sizeof(sharedmem), PROT_READ | PROT_WRITE, MAP_SHARED, fd, 0);
```

for consistency, since I prefer explicitly casting the void pointers, it should be:

```
shmp = (sharedmem*) mmap(NULL,
            sizeof(sharedmem), PROT_READ | PROT_WRITE, MAP_SHARED, fd, 0);
```

---

### Page 633

Though it is not incorrect, the code after the `while` loop cannot be executed given the current form of the loop. The call to `fatal_error()` could be preceded by a call to `mq_close()` and the code after the loop could be deleted.

---

### Page 639

The macro definition in *ulogger.c*

```
#define SIGUPDATE SIGRTMIN+1
```

is correct as is, but in general, it's best to put parentheses around macro values that contain operators, i.e., `(SIGRTMIN+1)`.

---

### Page 642

The call to `cleanup()` at the end of *ulogger.c* is never reached and unnecessary.

---

### Page 663

The POSIX specification says that writes are atomic if the number of bytes is less than or equal to `PIPE_BUF`.

Writes are atomic **as long as** the number of bytes is smaller than `PIPE_BUF`.

should read:

Writes are atomic **if** the number of bytes is smaller than `PIPE_BUF`.

---

### Page 700

It is not an error, but it is inconsistent that I use `O_NDELAY` in this program instead of `O_NONBLOCK`. They are interchangeable:

```
while ( ((rawtext_fd = open(msg.raw_text_fifo,
            O_WRONLY | O_NDELAY)) == -1) && (tries < MAXTRIES) ) {
```

---

### Page 720

At the bottom of this page, in the short listing that suggests two ways to pass data to threads, the declarations `int counter = 0;` and `int *on_heap = (int*) malloc(sizeof int);` are file-scoped. This code is really a straw man that leads to the better solution, and cannot be compiled because `on_heap` is assigned the return value of a function call.

The `on_heap` variable should be initialized in the `main` program before any threads are created:

```
int *on_heap; /* Initialize in main() with on_heap = (int*) malloc(sizeof int); */
```

---

### Page 731

Not all of the output of `ps -L -eopid,ppid,tid,cmd | grep pthread_signal_demo` is shown in the listing.

---

### Page 793

The text should note that the program in Listing 17-1 doesn't reset standard input to blocking mode just before it terminates because the shell does this automatically. This is an important observation—it means that you cannot stop and restart this program or any program that enables non-blocking input unless you write explicit handlers for the SIGSTOP, SIGTSTP, and SIGCONT signals that reenables non-blocking input when the program is resumed.

---

### Page 835

Although the code is correct, the following changes make it more consistent:

```
if ( -1 == write(STDOUT_FILENO, prompt, strlen(prompt)) )
    fatal_error(errno, "write");
while ( read(STDIN_FILENO, &inbuf, 1) > 0 ) {
    inbuf = toupper(inbuf);
    if ( -1 == write(STDOUT_FILENO, &inbuf, 1) )
        fatal_error(errno, "write");
}
```

### Page 873

On Linux kernels since Linux 6.2, this program may not work unless you set the CAP\_SYS\_ADMIN capability on the binary. You'll need superuser privilege. After building the executable, add the line

```
$ sudo setcap cap_sys_admin=eip fakeinput_demo
```

before running:

```
$ ./fakeinput_demo
$ /usr/bin/echo 'This is a dangerous thing to do!'
This is a dangerous thing to do!
$
```