

Linux Kernel Development

Copyright © 2004 by Sams Publishing

International Standard Book Number: 0-672-32512-8

Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an "as is" basis. The author and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book.

When reviewing corrections, always check the print number of your book. Corrections are made to printed books with each subsequent printing. To determine the printing of your book, view the copyright page. The print number is right-most number on the line below the "First Printing" line. For example, the following indicates the 4th printing of a title.

First Printing: August 2003

06 05 04 03 10 9 8 7 6 5 4

Misprint	Correction																
<p>Page 12, first paragraph</p> <p>The kernel stack is neither large nor dynamic; it is small and fixed in size. The kernel stack is fixed at 4 KB on 32-bit architectures and 8 KB on 64-bit architectures.</p>	<p>The kernel stack is neither large nor dynamic; it is small and fixed in size. The kernel stack is fixed at 8 KB on 32-bit architectures and 16 KB on most 64-bit architectures.</p>																
<p>Page 34, Figure 3.1</p> <p>Bottom row of figure</p> <table border="0" data-bbox="176 634 1047 732"> <thead> <tr> <th>Minimum</th> <th>Default</th> <th>Maximum</th> </tr> </thead> <tbody> <tr> <td>10 ms</td> <td>150ms</td> <td>300ms</td> </tr> </tbody> </table>	Minimum	Default	Maximum	10 ms	150ms	300ms	<table border="0" data-bbox="1047 586 1938 732"> <thead> <tr> <th>Minimum</th> <th>Default</th> <th>Maximum</th> </tr> </thead> <tbody> <tr> <td>10 ms</td> <td>100ms</td> <td>200ms</td> </tr> </tbody> </table>	Minimum	Default	Maximum	10 ms	100ms	200ms				
Minimum	Default	Maximum															
10 ms	150ms	300ms															
Minimum	Default	Maximum															
10 ms	100ms	200ms															
<p>Page 39, code snippet at bottom of page</p> <pre>struct prio_array array = rq->active;</pre>	<pre>struct prio_array*array = rq->active;</pre>																
<p>Page 42, Table 3.1, last line</p> <table border="0" data-bbox="176 878 1047 927"> <thead> <tr> <th>Maximum</th> <th>3200ms</th> <th>high</th> <th>low</th> </tr> </thead> <tbody> <tr> <td></td> <td></td> <td></td> <td></td> </tr> </tbody> </table>	Maximum	3200ms	high	low					<table border="0" data-bbox="1047 878 1938 927"> <thead> <tr> <th>Maximum</th> <th>200ms</th> <th>high</th> <th>low</th> </tr> </thead> <tbody> <tr> <td></td> <td></td> <td></td> <td></td> </tr> </tbody> </table>	Maximum	200ms	high	low				
Maximum	3200ms	high	low														
Maximum	200ms	high	low														
<p>Page 44, first paragraph</p> <p>Wait queues are created statically via DECLARE_WAIT_QUEUE_HEAD() or dynamically via <code>init_waitqueue_head()</code>.</p>	<p>Wait queues are created statically via DECLARE_WAITQUEUE() or dynamically via <code>init_waitqueue_head()</code>.</p>																
<p>Page 48, last paragraph</p> <p>The first change in supporting kernel preemption was the addition of a preemption counter, <code>preempt_count</code>, to each process's task_struct. This counter begins...</p>	<p>The first change in supporting kernel preemption was the addition of a preemption counter, <code>preempt_count</code>, to each process's thread_info structure. This counter begins...</p>																

<p>page 49, first line after the "Real-Time" heading</p> <p>Linux provides two real-time scheduling policies, SCHED_FF and SCHED_RR.</p>	<p>Linux provides two real-time scheduling policies, SCHED_FIFO and SCHED_RR.</p>
<p>Page 50, Table 3.3: The descriptions for <code>sched_setaffinity()</code> and <code>sched_getaffinity()</code> are swapped.</p>	
<p>Page 62, code snippet at bottom of page</p> <pre>#define NR_open 5</pre>	<pre>#define_NR_open 5</pre>
<p>Page 63, last bullet on page</p> <p>Chapter 13, "Virtual Filesystems," provides more details.</p>	<p>Chapter 11, "Virtual Filesystems," provides more details.</p>
<p>Page 67, text in "Top Halves Versus Bottom Halves" section</p> <p>The bottom half runs later, at a more convenient time, with all interrupts disabled..</p>	<p>The bottom half runs later, at a more convenient time, with all interrupts enabled..</p>
<p>Page 75, first paragraph</p> <p>If so, it calls <code>hardware_irq_event()</code> to run the installed interrupt handlers for the line.</p>	<p>If so, it calls <code>handle_irq_event()</code> to run the installed interrupt handlers for the line.</p>
<p>Page 77, second paragraph in "Interrupt Control" section</p> <p>Neither disabling interrupt deliver nor disabling kernel preemption provides any protection from concurrent access from another processor, however.</p>	<p>Neither disabling interrupt delivery nor disabling kernel preemption provides any protection from concurrent access from another processor, however.</p>

<p>Page 79, paragraph before "Status of the Interrupt System"</p> <p>Disabling the line disables interrupt deliver for <i>all</i> devices on the line.</p>	<p>Disabling the line disables interrupt delivery for <i>all</i> devices on the line.</p>
<p>Page 80, all occurrences of deliver in Table 5.1 should be delivery</p>	
<p>Page 83, fourth paragraph</p> <p>The top half could mark whether the bottom half would run by sitting a bit in a 32-bit integer.</p>	<p>The top half could mark whether the bottom half would run by setting a bit in a 32-bit integer.</p>
<p>Page 89, last paragraph</p> <p>Tasklets are represented by the tasklist_struct structure.</p>	<p>Tasklets are represented by the tasklet_struct structure.</p>
<p>Page 90, first paragraph in "Scheduling Tasklets" section</p> <p><i>Scheduled</i> tasklets (the equivalent of raised softirqs) are stored in two per-processor structures: tasklist_vec (for regular tasklets) and tasklet_hi_vec (for high-priority tasklets).</p>	<p><i>Scheduled</i> tasklets (the equivalent of raised softirqs) are stored in two per-processor structures: tasklet_vec (for regular tasklets) and tasklet_hi_vec (for high-priority tasklets).</p>
<p>Page 90, third bullet</p> <p>Add the tasklet to-be-scheduled to the head of the tasklet_vec or tasklist_hi_vec linked list, which is unique to each processor in the system.</p>	<p>Add the tasklet to-be-scheduled to the head of the tasklet_vec or tasklet_hi_vec linked list, which is unique to each processor in the system.</p>
<p>Page 92, second paragraph</p> <p>Both of these macros statically create a struct tasklist_struct with the given name.</p>	<p>Both of these macros statically create a struct tasklet_struct with the given name.</p>

<p>Page 99, last paragraph</p> <p>To create the structure statically at run-time:</p>	<p>To create the structure statically at compile-time:</p>
<p>page 125, second paragraph below the "Spin Locks and Bottom Halves" heading</p> <p>Because a bottom half may preempt process context code, if data is shared between a bottom half process context, you must protect the data in process context with both a lock and the disabling of bottom halves.</p>	<p>Because a bottom half may preempt process context code, if data is shared between a bottom half and process context, you must protect the data in process context with both a lock and the disabling of bottom halves.</p>
<p>Page 128, first paragraph</p> <p>...one of the tasks on the wait queue will be awakened up so that it can acquire the semaphore.</p>	<p>...one of the tasks on the wait queue will be woken up so that it can acquire the semaphore.</p>
<p>Page 128, middle paragraph, next-to-last sentence</p> <p>Additionally, unlike spin locks, semaphores do not disable kernel preemption and, consequently, code holding a spin lock can be preempted.</p>	<p>Additionally, unlike spin locks, semaphores do not disable kernel preemption and, consequently, code holding a semaphore can be preempted.</p>
<p>Page 132, paragraph before Table 8.7</p> <p>After the event has occurred, calling complete() signals all waiting tasks to wake up.</p>	<p>After the event has occurred, calling complete() signals a waiting task to wake up.</p>
<p>Page 133, Table 8.7</p> <p>Signals any waiting tasks to wake up</p>	<p>Signals a waiting task to wake up</p>
<p>Page 153, paragraph in middle of page</p> <p>The xtime.v_nsec value stores the number of nanoseconds that have elapsed in the last second.</p>	<p>The xtime.tv_nsec value stores the number of nanoseconds that have elapsed in the last second.</p>

<p>Page 165, paragraph after second bulleted list</p> <p>The actual use and layout of the memory zones is architecture independent.</p>	<p>The actual use and layout of the memory zones is architecture dependent.</p>
<p>Page 172, Table 10.5, entry for GFP_HIGHUSER</p> <p>This is an allocation from ZONE_HIGHMEM and might block.</p>	<p>This is an allocation from ZONE_HIGHMEM and might block.</p>
<p>Page 173, second paragraph</p> <p>On the far other end of the spectrum is the GFP_ATOMIC flag.</p>	<p>On the far other end of the spectrum is the GTP_ATOMIC flag.</p>
<p>page 181, first line on the page</p> <p>This creates a cache named task_struct, which stores objects of type struct task_struct.</p>	<p>This creates a cache named task_struct_cachep, which stores objects of type struct task_struct.</p>
<p>Page 194, third line</p> <p>/* file creation timestamp */</p>	<p>/* inode change time */</p>
<p>Page 208, second paragraph in "Data Structures Associated with a Process" section</p> <p>The address of this table is pointed to by the files entry in the processor descriptor.</p>	<p>The address of this table is pointed to by the files entry in the process descriptor.</p>
<p>Page 215, paragraph before "The bio structure"</p> <p>...into many multiple buffer_head structures.</p>	<p>...into multiple buffer_head structures.</p>
<p>Page 217, first paragraph after first code snippet</p> <p>In each given block I/O operation, there are bi_vcnt vectors in the bio_vec array starting with bi_io_vecs.</p>	<p>In each given block I/O operation, there are bi_vcnt vectors in the bio_vec array starting with bi_io_vec.</p>

<p>Page 217, paragraph in middle of page</p> <p>Table 12.3 is a diagram of the relationship between the <code>bio</code> structure, the <code>bio_vec</code> structure, and the <code>page</code> structure.</p>	<p>Figure 12.2 is a diagram of the relationship between the <code>bio</code> structure, the <code>bio_vec</code> structure, and the <code>page</code> structure.</p>
<p>page 226, footnote 3 at the bottom of the page</p> <p>Newer versions of glibc implement <code>malloc()</code> via <code>mmap()</code> and not <code>brk()</code>.</p>	<p>Newer versions of glibc implement <code>malloc()</code> via <code>mmap()</code> and <code>brk()</code>.</p>
<p>Page 230, next-to-last paragraph</p> <p>Thus, <code>vm_end - vm_start</code> is the size (length) in bytes of the interval.</p>	<p>Thus, <code>vm_end - vm_start+1</code> is the size (length) in bytes of the interval.</p>
<p>page 231, 8th entry in Table 13.1</p> <p><code>VM_MAYSHARE</code> The <code>VM_SHARE</code> flag can be set</p>	<p><code>VM_MAYSHARE</code> The <code>VM_SHARED</code> flag can be set</p>
<p>Page 244, first paragraph in "The <code>address_space</code> Object" section</p> <p>Checking the page cache to see if certain data has been cached is rendered more difficult because of the noncontiguous nature of the blocks that can up each page.</p>	<p>Checking the page cache to see if certain data has been cached is rendered more difficult because of the noncontiguous nature of the blocks that can make up each page.</p>
<p>page 245, second to last line of code</p> <pre>struct address_space *assoc_mapping; /* associated buffres */</pre>	<pre>struct address_space *assoc_mapping; /* associated buffers */</pre>
<p>Page 257, first paragraph</p> <p>The lone disadvantage of a circular buffer—the possibility of loosing messages—is a small price to pay for the simplicity and robustness it affords.</p>	<p>The lone disadvantage of a circular buffer—the possibility of losing messages—is a small price to pay for the simplicity and robustness it affords.</p>

<p>Page 276, first paragraph in "Byte Order" section</p> <p>The byte ordering is called <i>big-endian</i> if the most significant byte is encoding first with the remaining bytes decreasing in significance.</p>	<p>The byte ordering is called <i>big-endian</i> if the most significant byte is encoded first with the remaining bytes decreasing in significance.</p>
<p>Page 278, first line of code</p> <p>u23 __cpu_to_be32(u32);</p>	<p>u32 __cpu_to_be32(u32);</p>

This errata sheet is intended to provide updated technical information. Spelling and grammar misprints are updated during the reprint process, but are not listed on this errata sheet.