

# Contents

<b>Preface</b>	<b>iii</b>
<b>1 Hardware Basics</b>	<b>1</b>
1.1 The CPU . . . . .	2
1.2 Memory . . . . .	4
1.3 Buses . . . . .	4
1.4 Controllers and Peripherals . . . . .	5
1.5 Address Spaces . . . . .	5
1.6 Timers . . . . .	6
<b>2 Software Basics</b>	<b>7</b>
2.1 Computer Languages . . . . .	7
2.1.1 Assembly Languages . . . . .	7
2.1.2 The C Programming Language and Compiler . . . . .	8
2.1.3 Linkers . . . . .	9
2.2 What is an Operating System? . . . . .	9
2.2.1 Memory management . . . . .	10
2.2.2 Processes . . . . .	10
2.2.3 Device drivers . . . . .	11
2.2.4 The Filesystems . . . . .	11
2.3 Kernel Data Structures . . . . .	11
2.3.1 Linked Lists . . . . .	12
2.3.2 Hash Tables . . . . .	12
2.3.3 Abstract Interfaces . . . . .	13
<b>3 Memory Management</b>	<b>15</b>
3.1 An Abstract Model of Virtual Memory . . . . .	16
3.1.1 Demand Paging . . . . .	18
3.1.2 Swapping . . . . .	19
3.1.3 Shared Virtual Memory . . . . .	19
3.1.4 Physical and Virtual Addressing Modes . . . . .	19
3.1.5 Access Control . . . . .	20

---

3.2	Caches . . . . .	21
3.3	Linux Page Tables . . . . .	22
3.4	Page Allocation and Deallocation . . . . .	23
3.4.1	Page Allocation . . . . .	24
3.4.2	Page Deallocation . . . . .	24
3.5	Memory Mapping . . . . .	25
3.6	Demand Paging . . . . .	26
3.7	The Linux Page Cache . . . . .	27
3.8	Swapping Out and Discarding Pages . . . . .	28
3.8.1	Reducing the Size of the Page and Buffer Caches . . . . .	29
3.8.2	Swapping Out System V Shared Memory Pages . . . . .	30
3.8.3	Swapping Out and Discarding Pages . . . . .	30
3.9	The Swap Cache . . . . .	31
3.10	Swapping Pages In . . . . .	32
<b>4</b>	<b>Processes</b>	<b>35</b>
4.1	Linux Processes . . . . .	36
4.2	Identifiers . . . . .	38
4.3	Scheduling . . . . .	39
4.3.1	Scheduling in Multiprocessor Systems . . . . .	41
4.4	Files . . . . .	42
4.5	Virtual Memory . . . . .	43
4.6	Creating a Process . . . . .	45
4.7	Times and Timers . . . . .	46
4.8	Executing Programs . . . . .	47
4.8.1	ELF . . . . .	48
4.8.2	Script Files . . . . .	50
<b>5</b>	<b>Interprocess Communication Mechanisms</b>	<b>51</b>
5.1	Signals . . . . .	51
5.2	Pipes . . . . .	53
5.3	Sockets . . . . .	55
5.3.1	System V IPC Mechanisms . . . . .	55
5.3.2	Message Queues . . . . .	55
5.3.3	Semaphores . . . . .	56
5.3.4	Shared Memory . . . . .	58
<b>6</b>	<b>PCI</b>	<b>61</b>
6.1	PCI Address Spaces . . . . .	61
6.2	PCI Configuration Headers . . . . .	62
6.3	PCI I/O and PCI Memory Addresses . . . . .	64

---

6.4	PCI-ISA Bridges . . . . .	64
6.5	PCI-PCI Bridges . . . . .	65
6.5.1	PCI-PCI Bridges: PCI I/O and PCI Memory Windows . . . . .	65
6.5.2	PCI-PCI Bridges: PCI Configuration Cycles and PCI Bus Numbering . . . . .	65
6.6	Linux PCI Initialization . . . . .	66
6.6.1	The Linux Kernel PCI Data Structures . . . . .	67
6.6.2	The PCI Device Driver . . . . .	68
6.6.3	PCI BIOS Functions . . . . .	70
6.6.4	PCI Fixup . . . . .	72
<b>7</b>	<b>Interrupts and Interrupt Handling</b>	<b>75</b>
7.1	Programmable Interrupt Controllers . . . . .	77
7.2	Initializing the Interrupt Handling Data Structures . . . . .	77
7.3	Interrupt Handling . . . . .	78
<b>8</b>	<b>Device Drivers</b>	<b>81</b>
8.1	Polling and Interrupts . . . . .	82
8.2	Direct Memory Access (DMA) . . . . .	84
8.3	Memory . . . . .	85
8.4	Interfacing Device Drivers with the Kernel . . . . .	85
8.4.1	Character Devices . . . . .	86
8.4.2	Block Devices . . . . .	87
8.5	Hard Disks . . . . .	88
8.5.1	IDE Disks . . . . .	90
8.5.2	Initializing the IDE Subsystem . . . . .	91
8.5.3	SCSI Disks . . . . .	91
8.6	Network Devices . . . . .	95
8.6.1	Initializing Network Devices . . . . .	96
<b>9</b>	<b>The File system</b>	<b>99</b>
9.1	The Second Extended File system (EXT2) . . . . .	101
9.1.1	The EXT2 Inode . . . . .	102
9.1.2	The EXT2 Superblock . . . . .	103
9.1.3	The EXT2 Group Descriptor . . . . .	104
9.1.4	EXT2 Directories . . . . .	104
9.1.5	Finding a File in an EXT2 File System . . . . .	105
9.1.6	Changing the Size of a File in an EXT2 File System . . . . .	106
9.2	The Virtual File System (VFS) . . . . .	107
9.2.1	The VFS Superblock . . . . .	109
9.2.2	The VFS Inode . . . . .	109

---

9.2.3	Registering the File Systems . . . . .	110
9.2.4	Mounting a File System . . . . .	110
9.2.5	Finding a File in the Virtual File System . . . . .	112
9.2.6	Creating a File in the Virtual File System . . . . .	112
9.2.7	Unmounting a File System . . . . .	112
9.2.8	The VFS Inode Cache . . . . .	113
9.2.9	The Directory Cache . . . . .	113
9.3	The Buffer Cache . . . . .	114
9.3.1	The <code>bdfflush</code> Kernel Daemon . . . . .	116
9.3.2	The <code>update</code> Process . . . . .	116
9.4	The <code>/proc</code> File System . . . . .	116
9.5	Device Special Files . . . . .	117
<b>10</b>	<b>Networks</b>	<b>119</b>
10.1	An Overview of TCP/IP Networking . . . . .	119
10.2	The Linux TCP/IP Networking Layers . . . . .	122
10.3	The BSD Socket Interface . . . . .	124
10.4	The INET Socket Layer . . . . .	125
10.4.1	Creating a BSD Socket . . . . .	127
10.4.2	Binding an Address to an INET BSD Socket . . . . .	127
10.4.3	Making a Connection on an INET BSD Socket . . . . .	128
10.4.4	Listening on an INET BSD Socket . . . . .	129
10.4.5	Accepting Connection Requests . . . . .	130
10.5	The IP Layer . . . . .	130
10.5.1	Socket Buffers . . . . .	130
10.5.2	Receiving IP Packets . . . . .	131
10.5.3	Sending IP Packets . . . . .	132
10.5.4	Data Fragmentation . . . . .	133
10.6	The Address Resolution Protocol (ARP) . . . . .	133
10.7	IP Routing . . . . .	135
10.7.1	The Route Cache . . . . .	135
10.7.2	The Forwarding Information Database . . . . .	136
<b>11</b>	<b>Kernel Mechanisms</b>	<b>139</b>
11.1	Bottom Half Handling . . . . .	139
11.2	Task Queues . . . . .	140
11.3	Timers . . . . .	141
11.4	Wait Queues . . . . .	142
11.5	Buzz Locks . . . . .	143
11.6	Semaphores . . . . .	143

---

<b>12 Modules</b>	<b>145</b>
12.1 Loading a Module . . . . .	146
12.2 Unloading a Module . . . . .	148
<b>13 Processors</b>	<b>151</b>
13.1 X86 . . . . .	151
13.2 ARM . . . . .	151
13.3 Alpha AXP Processor . . . . .	152
<b>14 The Linux Kernel Sources</b>	<b>153</b>
<b>A Linux Data Structures</b>	<b>159</b>
<b>B Useful Web and FTP Sites</b>	<b>177</b>
<b>C Linux Documentation Project Manifesto</b>	<b>179</b>
C.1 Overview . . . . .	179
C.2 Getting Involved . . . . .	180
C.3 Current Projects . . . . .	180
C.4 FTP sites for LDP works . . . . .	180
C.5 Documentation Conventions . . . . .	180
C.6 Copyright and License . . . . .	181
C.7 Publishing LDP Manuals . . . . .	181
<b>D The GNU General Public License</b>	<b>183</b>
D.1 Preamble . . . . .	183
D.2 Terms and Conditions . . . . .	184
D.3 How to Apply These Terms . . . . .	188
<b>Glossary</b>	<b>191</b>
<b>Bibliography</b>	<b>194</b>