



Index

SYMBOLS

./configure command, Valgrind installation, 99
 /boot directory layout, 188-189
 /etc/syslog.conf configuration file, 213
 /proc file system, 112
 administrative applications, 113
 entries, 114-116
 interfaces, 112
 relationship with sysctl functions, 112
 utilities
 /proc/vmstat file, 146
 /sys/kernel file, 140-143
 /sys/vm file, 144-145
 buddyinfo file, 117
 bus directory, 118
 cmdline file, 121
 cpuinfo file, 122
 dev/fs file, 138
 diskstats file, 122-126
 free file, 134
 interrupts file, 127
 kallsyms file, 128
 kcore file, 128
 kmsg file, 128
 loadavg file, 129
 lsdev file, 128
 meminfo file, 129
 mpstat, 134
 net directory, 135-136
 slabinfo file, 136
 stat file, 137
 sys directory, 138
 sysrq-trigger file, 145
 vmstat, 131-133
 Web resources, 149
 /proc/mm support option (UML configuration), 298
 /proc/vmstat file, 146
 /sys/kernel file, 140-143
 /sys/vm file, 144-145
 -l hostlist option (syslogd daemon), 213
 1000 memory overrun, 97
 2.5.1 pre5, kernel mailing list Oops, 204-207

2G/2G host address space split option (UML configuration), 297
 3c59x network driver, kernel mailing list Oops, 200-204

A

a option
 Account Time Tracing, 253
 unlimit gdb command, 63
 acquisition of data, LTT (Linux Trace Toolkit), 243-244
 action field (syslog.conf configuration file), 215
 activation
 kdb debugger, 352
 serial console port, 352-353
 addr field (register_probe interface), 323
 Address Resolution Protocol (ARP cache), 178
 ARP table, 135
 manipulation, 179
 administrative applications (/proc file system), 113
 Akkerman, Wichert, 201
 aligning structures, performance counters, 279
 Alt-SysRq-command, 145
 an option (netstat command), 181
 analysis
 Oops messages, 190-195, 372
 tracing processes, LTT (Linux Trace Toolkit), 245
 applications
 debugging
 Dprobes, 316-331
 gdb, 54-76
 memory management debugging, 82-109
 symbols that increase executable size, 77
 techniques, 78-79
 Kprobes, 330-331
 trace, 234
 UML (User-Mode Linux), 292
 architecture component tasks, LTT (Linux Trace Toolkit), 235
 ARP (Address Resolution Protocol) cache, 178
 ARP table, 135
 manipulation, 179
 arp command, 179

arrays, 280
 assembly code, generation of by gcc `-s` option,
 199-200
 attach command, 56

B

b option
 gprof profiler, 17
 sar utility, 126

back traces
 magic key sequence, 174-176
 PID 2, 362

backtrace command, 60

binary interface, 112

bio_alloc() routine, 205

blocking technique, performance counters, 281,
 285-287

bogomips entry (cpuinfo file, /proc file system),
 122

booting UML (User-Mode Linux), 307

branch coverage, 40

break 30 command, 46

break 36 command, 48

break command, 60

break points, 54

BSQ_CACHE_REFERENCE counter, 283

bt command, 355, 397
 back traces, 396
 viewing the back trace for process 1960, 399

bta (Back Trace for All Processes) kdb command,
 362

buddyinfo file, 117

buffer overruns, 82

building the kernel
 crash dumps, 376-378
 kdb debugger, 350-351
 kgdb debugger, 339-341
 Kprobes, 320-321
 LTT (Linux Trace Toolkit), 240-242
 Oops message analysis, 193-195
 UML (User-Mode Linux), 296-304
configuration options, 297-299
gcov support, 302
gprof support, 303
Kernel backing menu, 301
transports, 299-301
UML Network Devices menu, 299

Built-in Kernel Debugger support option, 335

bunzip2 command, 98, 336

bus directory (/proc file system), 118

bus error messages, 79

busy waiting, 202

bzip2 command, 348

C

c option (CPU ID tracing)
 tracevisualizer program, 253
 unlimit gdb command, 63

cache misses
 data cache misses, 282
 performance counters, 278

cache profiling (Valgrind), 106-109

call graphs, 14

call stacks, 55

called functions (kprof), 31

calloc() function, memory management, 83

cg_annotate command, 108

character-based interface (/proc file system), 112

checkers (memory-management)
 Electric Fence package, 93-96
 MEMWATCH, 84-86
 Valgrind, 97
cache profiling, 106-109
finding memory errors, 101-106
installation, 98-101
limitations, 109
 Web resources, 81, 109-111
 YAMD package, 86-93

clear_inode method (VFS), 359

client requirements, Netdump, 391-392

clock() function, execution time, 6, 9-10

cmdline file, 121

code coverage analysis, 38
 gcov
branch coverage, 40
gcc options, 41-42, 51
loop coverage, 40
 logic errors, 39-40

combined loops (performance counters), 281

commands
 ./configure, 99
 arp, 179
 bt, 355, 396-399
 bunzip2, 98, 336
 bzip2, 348
 cg_annotate, 108
 cont, 344
 cp, 189
 cp /usr/src/linux/linux /usr/src/uml/linux, 304
 crash dump interactive analysis, 373-374
 date, 4
 diff, 202
 dis, 403
 echo, 152
 evlconfig, 229
 evlgentmpls, 229

- evlogmgr, 228-229
- evlrc, 229
- files, 404-405
- gdb, 56-60, 74
 - break* 30, 46
 - break* 46, 48
 - cont*, 48-49
 - list*, 46-48
 - print array*, 46
 - print array[0]*, 49
 - quit*, 48-49
 - run* 1, 46, 49
 - set array=0*, 46
 - set array[0]=0*, 49
 - step*, 48-49
- go, 356
- help, 410
- info breakpoints, 307, 346
- info gprof, 15
- info threads, 309
- kdb
 - bta* command, 362
 - dmseg* command, 361
 - ps* command, 361
- kdb debugger, 353-354
- kill, 245
- lcrash commands, 385-386
- linux, 294
- ln, 190
- log, 400-403
- logger (Syslog), 218-219
- logrotate, 221
- mach, 396
- make, 99
- make bzImage, 265, 320
- make clean, 265, 320
- make install, 88, 99
- make modules, 265, 321
- make modules_install, 265, 321
- make xconfig, 263, 320
- make xconfig ARCH=um, 299, 303
- man gprof, 15
- man ps, 154
- md, 366
- mount, 344, 399
- mount root_fs mnt -o loop, 304
- netstat, 180-182
- nm, 189
- objdump, 356
- opannotate, 273
- opcontrol --init, 282
- op_help, 282
- patch, 202, 296
- perror, 188, 208-209
- pgrep bash, 169
- ps, 116, 245, 397-398
- quit, 410
- rpm, 16
- run-yamd, 92
- runq, 403
- slog_fwd, 231
- stat, 387
- stty, 341
- sys, 395-396
- tar, 98
- task, 405, 409
- thread, 309
- thread apply, 309
- time, 5-6
- tracedaemon, 242
- tracevisualizer, 242, 251
- umount mnt, 304
- uname, 321
- uname -a, 265
- valgrind ls -all, 100
- compilation of programs (gdb), 59**
- compile-time flags (MEMWATCH), 85**
- config file, 189**
- configuration**
 - Diskdump, 393
 - kernel configuration, crash dumps, 374-376
 - network interface, ifconfig utility, 178
 - UML (User-Mode Linux), 297-299
- CONFIG_DEBUGSYM debugging symbol, 302**
- CONFIG_PT_PROXY debugging symbol, 302**
- consoles (UML), 310**
- consumer interface, 227**
- cont array command, 48**
- cont command, 49, 60, 344**
- copy-on-write (COW) file system files, 293**
- core files, debugging applications with gdb, 63-65**
- core-file command, 64**
- COW (copy-on-write) file system files, 293**
- Cox, Alan, 201**
- cp /usr/src/linux/linux /usr/src/uml/linux**
 - command, 304
- cp command, 189**
- cpuinfo file, 122**
- CPU_CLOCK_UNHALTED counter, 281**
- crash dumps, 372**
 - Diskdump, 392-393
 - interactive analysis, 373-374
 - kernel configuration, 374-376
 - lcrash commands, 385-386
 - mcore crash dump, 393-405, 410
 - Netdump, 391-392
 - patching and building the kernel, 376-378
 - report generation, 373

- sample system panic and crash dump taken, 386, 390
- steps, 379-384
- Web resources, 411
- crash tool, 373
- CRASH_DUMP option (LKCD kernel crash dump), 375
- CRASH_DUMP_BLOCKDEV option (LKCD kernel crash dump), 375
- CRASH_DUMP_COMPRESS_GZIP option (LKCD kernel crash dump), 376
- CRASH_DUMP_COMPRESS_RLE option (LKCD kernel crash dump), 376
- CRASH_DUMP_MEMDEV option (LKCD kernel crash dump), 375
- CRASH_DUMP_NETDEV option (LKCD kernel crash dump), 375
- CRASH_DUMP_SOFTBOOT option (LKCD kernel crash dump), 375
- cron job, 228

D

Daemon transport, UML packet exchange, 301

data (LTT)

- acquisition, 243-244
- interpretation, 248-253
- recording, 244
- reviewing, 256-259
- saving, oprofile, 277

data cache misses, 282

Data Display Debugger. See ddd

date command, 4

dbench program execution time, 4

ddd (Data Display Debugger), 54

- features, 66
- installation, 66-67
- invocation, 67

dead code, 38-39

Debian 3.0 root image, 305

debuggers, 334-335

- kdb, 348
 - activation, 352-353
 - building the kernel, 350-351
 - commands, 353-354
 - debug sessions, 354-368
 - kernel patching, 348-350
 - Web resources, 369

kgdb, 335

- building the kernel, 339-341
- checking the null modem serial cable setup, 341-343
- gdb sessions, 343
- gdb thread analysis, 347
- kernel patching, 336-338

- stopping kernel execution, 344-346

- Web resources, 369

debugging applications

Dprobes, 316

- characteristics, 318

- DPEH (Dprobes Event Handler), 317

- general steps, 318-319

- Kprobes, 319-331

- Probe Manager, 316

- Web resources, 331

gdb, 54

- commands, 56-59

- compiling program to be debugged, 59

- core files, 63-65

- GUIs (graphical user interfaces), 54, 65-76

- installation, 55

- invocation, 60-63

- operations, 54-55

memory management, 82

- Electric Fence package, 93-96

- errors, 83-84

- libc (library calls), 82-83

- MEMWATCH, 84-86

- Valgrind, 97-109

- YAMD package, 86, 91-93

printf() statements, 54

symbols, 77

Syslog, 222-226

techniques, 78-79

delete command (gdb), 60

dev/fs file (/proc file system), 138

development

- machines, kgdb requirement, 336

- UML file systems, 309

diff command, 202

Direct Rendering Manager (DRM), 357

dis command, 403

disk space (Syslog), 219-220

Diskdump, 372, 392-393

diskstats file, 122-126

- iostat utility, 122-125

- information reported, 123-124

- measurements of tuning results, 125-126

- sample output, 124

- x option, 124

- sar utility, 126

dlopen() function, 156

dlssystem() function, 156

dmap_pmap_verify() routine, 108

DMEMWATCH compile-time flag, 85

dmesg kdb command, 361

DMW_STUDIO compile-time flag, 85

dpcc, 330

DPEH (Dprobes Event Handler), 317

Dprobes (dynamic probes), 316

- characteristics, 318
- DPEH (Dprobes Event Handler), 317
- general steps, 318-319
- Kprobes, 319
 - applications, 330-331*
 - building the kernel, 320-321*
 - creating at an offset in a routine, 328-329*
 - finding active Kprobes, 328*
 - interfaces, 321-322*
 - Jprobes, 329-330*
 - makefile for a basic kernel module, 326*
 - registering and unregistering, 322-324*
 - sys_open sample, 324-326*
- Probe Manager, 316
- Web resources, 331

DRM (Direct Rendering Manager), 357

dumping traffic on a network (tcpdump), 179

dynamic probes. See Dprobes

E

echo command, 152

EIP (Extended Instruction Pointer) register, 354

Eldredge, Nate, YAMD package, 86

Electric Fence package, 93-96

entries (/proc file system), 114-116

error messages, 188

- config files, 189
- event logging, 212-213, 226
 - evlog packages, 231*
 - evlogd daemon, 226*
 - evlogmgr command, 228*
 - evlogrmt daemon, 230-231*
 - features, 227*
 - forwarding Syslog messages to event log, 231*
 - interfaces, 227*
 - log management, 228*
 - log records, 228*
 - remote logging, 229-230*
 - utilities, 228-229*
 - Web resources, 232*
- kernel files, 188-189

Oops messages

- analysis, 190-195*
- applying gdb to display jfs_mount, 199*
- generation of assembly code by gcc -s option, 199-200*
- kernel mailing lists, 200-207*
- processing using ksymoops, 196-198*
- Web resources, 209*

perror command, 208-209

Syslog, 212-226

system.map, 189-190

Error Record Template Repository, 226

Ethereal, browsing network traffic, 180

Ethertap transport, UML packet exchange, 300

Event Consolidation Host, 229

Event Graph (tracevisualizer program), 249

Event Handler (Dprobes). See DPEH

event logging, 212-213, 226

- evlog packages, 231
- evlogd daemon, 226
- evlogmgr command, 228
- evlogrmt daemon, 230-231
- features, 227
- forwarding Syslog messages to event log, 231
- interfaces, 227
- log management, 228
- log records, 228
- remote logging, 229-230
- utilities, 228-229
- Web resources, 232

evlconfig command, 229

evlfacility utility, 229

evlgentmpls command, 229

evlnotify utility, 229

evlog packages, 231

evlogd daemon, 226

evlogmgr command, 228-229

evlogrmt daemon, 230-231

evlquery utility, 229

evlvc command, 229

evlview utility, 229

executable size (programs), 77

execution time, 2

- clock() function, 6, 9-10
- date command, 4
- gettimeofday() function, 11-13
- gprof profiler, 13-14
 - b option, 17*
 - field descriptions, 28-31*
 - for loops/if statements, 20, 26*
 - fprofile-arcs option, 18*
 - fiest-coverage option, 18*
 - kprof, 31-32*
 - pg gcc option, 15-16*
- stopwatch method, 3
- time command, 5-6

execution traces, 234

applications, 234

LTT (Linux Trace Toolkit), 234

analyzing tracing processes, 245

architecture component tasks, 235

building the kernel, 240-242

building user-level tools, 242

data acquisition, 243-244

data recording, 244

data reviewing text tools, 256-259
interpretation of data, 248-253
package and installation, 236-240
tarring a subdirectory sample, 253-256
Web resources, 259

Extended Instruction Pointer (EIP) register, 354

F

f option (Event Types Tracing), tracevisualizer program, 252

facility code, 215

fault_handler function, 323

features

ddd (Data Display Debugger), 66
 event logging, 227
 Insight, 70-71

field descriptions (gprof profiler), 28-31

fields (/proc file system)

dev/fs file, 138-139
 mpstat utility, 134
 stat file, 137-138
 /sys/kernel file, 140-142
 /sys/vm file, 144

file systems, (UML), 305, 311

files command, 404-405

flat profile view, 31

flat profiles, 13-14

for loops, 20, 26

Forwarding Plug-ins, 229

forwarding Syslog messages to event log, 231

fprofile-arcs (gcc option), 41

fprofile-arcs option (gprof profiler), 18

frames (stack frames), 55

free files, 134

free() function, memory management, 83

fscklog_init() function, 101

ftest-coverage option (gcc), 42

ftest-coverage option (gprof profiler), 18

functions

clock(), 6, 9-10
 gettimeofday(), 11-13
 memory management, 83
 sqrt(), 20

G

g compiler option (gdb), 59

g option (Electric Fence package), 94

gcc options

gcov, 41-42, 51
 -fprofile-arcs option, 41
 -ftest-coverage option, 42
 -s option, generation of assembly code, 199-200
 gprof profiler, 15-16

gcov

code coverage analysis

branch coverage, 40

gcc options, 41-42, 51

loop coverage, 40

logic errors, 39-40

UML support, 302

gdb (GNU debugger), 54

commands, 46-48, 56-60, 74

compiling program to be debugged, 59

core files, 63-65

displaying jfs_mount code, 199

GUIs (graphical user interfaces), 54, 65

ddd, 66-67

Insight, 70-76

installation, 55

invocation, 60-63

operations, 54-55

sessions, 343

thread analysis, 347

UML sessions, 306-308

UML thread analysis, 308

Web resources, 53, 80

gdb executable_file corefile command (gdb), 56

gettimeofday() function, 11-13

GLOBAL_POWER_EVENTS counter, 281

GNU debugger. See gdb

go command, 356

gprof profiler

performance tuning, 13-14

-b option, 17

field descriptions, 28-31

for loops/if statements, 20, 26

fprofile-arcs option, 18

ftest-coverage option, 18

kprof, 31-32

pg gcc option, 15-16

UML support, 303

graph view (kprof), 31

graphical user interfaces. See GUIs

Graphviz, 32

GUIs (graphical user interfaces), 54

ddd, 54, 66-67

Insight, 54, 70-76

H

h option (syslogd daemon), 213

hardware counters, 277

aligning structures, 279

cache misses, loop grouping, 278-287

packing arrays, 280

padding structures, 279

hardware watch points, 54-55

help command, 410

hierarchical profile view (kprof), 31
 Highmem support option (UML configuration), 298
 HoneyPot ProcFS option (UML configuration), 297
 Host filesystem option (UML configuration), 297
 HUP signals, 213

I

ide_hwgroup_t structure, 365
 ide_timer_expiry routine, 368
 if statements, 20, 26
 ifconfig utility, 178
 info breakpoints command, 307, 346
 info gprof command, 15
 info threads command, 74, 309
 initial_processing() routine, 101
 Insight, 54
 features, 70-71
 installation, 71-76
installation
 ddd (Data Display Debugger), 66-67
 gdb, 55
 Insight, 71-76
 kprof, 32
 LTT (Linux Trace Toolkit), 236-240
 Valgrind, 98-101
 YAMD package, 88
instrumentation profilers, 262
 interactive analysis, crash dump commands, 373-374
interfaces
 event logging, 227
 Kprobes, 321-322
 /proc file system, 112
interpretation of data, LTT (Linux Trace Toolkit), 248-253
interrupts files, 127
invocation
 ddd (Data Display Debugger), 67
 gdb, 60-63
ioctl call, 173
iostat utility, 122-125
 information reported, 123-124
 measurements of tuning results, 125-126
 sample output, 124
 x option, 124
IRQ trace points, 235

J

JFS (Journaled File System), 188, 309
 jfs_mount code
 applying gdb to display jfs_mount, 199
 modifying code to create Oops messages, 191-192
 jfs_mount code, 191-192

Journaled File System. See JFS
Jprobes (jumper probes), 329-330

K

kallsyms files, 128
 kcore files, 128
 kdb debugger, 348
 activation, 352-353
 building the kernel, 350-351
 commands, 353-354
 debug sessions, 354-368
 kernel patching, 348-350
 Web resources, 369
KDB modules option, 335
KDB off by default option, 335
kernel
 building, 264-265
 configuration, crash dumps, 374-376
 files, 188-189
 LTT (Linux Trace Toolkit), 240-242
 mailing lists (Oops messages), 200
 2.5.1-pre5, 204-207
 3c59x network driver, 200-204
 UML, 292
 applications, 292
 booting, 307
 consoles and serial lines, 310
 file systems, 305, 309-311
 gdb sessions, 306-308
 gdb thread analysis, 308
 introduction to kernel and root file system, 293-295
 limitations, 293
 networking support, 310
 patching and building kernel, 296-304
 root images, 304, 311
 shutting down, 312
 utilities, 312-313
 Web resources, 314
Kernel Dynamic Probes. See Kprobes
Kernel hacking menu, 301, 334
Kernel support for ELF binaries option (UML configuration), 297
Kernel support for MISC binaries option (UML configuration), 297
kernel-level debuggers, 334-335
 kdb, 348
 activation, 352-353
 building the kernel, 350-351
 commands, 353-354
 debug sessions, 354-368
 kernel patching, 348-350
 Web resources, 369

- kgdb, 335
 - building the kernel, 339-341*
 - checking the null modem serial cable setup, 341-343*
 - gdb sessions, 343*
 - gdb thread analysis, 347*
 - kernel patching, 336-338*
 - stopping kernel execution, 344-346*
 - Web resources, 369*
- kgdb debugger, 335
 - building the kernel, 339-341
 - checking the null modem serial cable setup, 341-343
 - gdb sessions, 343
 - gdb thread analysis, 347
 - kernel patching, 336-338
 - stopping kernel execution, 344-346
 - Web resources, 369
- kgdb8250 option (kgdb kernel command-line option), 342
- KGDB: Console messages through gdb option, 335
- KGDB: Enable kernel asserts option, 335
- KGDB: Remote (serial) kernel debugging with gdb option, 335
- KGDB: Thread analysis option, 335
- kgdbwait option (kgdb kernel command-line option), 342
- kill command, 245
- klogd utility, 213, 217-218
- kmsg files, 128
- Kprobes (Kernel Dynamic Probes), 319
 - applications, 330-331
 - building the kernel, 320-321
 - creating at an offset in a routine, 328-329
 - finding active Kprobes, 328
 - interfaces, 321-322
 - Jprobes, 329-330
 - makefile for a basic kernel module, 326
 - registering and unregistering, 322-324
 - sys_open sample, 324-326
- kprof, 31-32
- ksymoops, 196-198
- L**
- lcrash commands, 385-386
- lcrash tool, 373
- lefence option (Electric Fence package), 94
- libc (library calls), 82
 - calloc() function, 83
 - free() function, 83
 - malloc() function, 83
 - realloc() function, 83
- library calls. *See* libc
- lilo boot loader, 321
- limitations
 - UML (User-Mode Linux), 293
 - Valgrind, 109
- Lindh, Johan, MEMWATCH package, 84
- linux command, 294
- Linux Kernel Crash Dump. *See* LKCD
- Linux Trace Toolkit. *See* LTT
- list command, 46-48, 60, 75
- list of callers (kprof), 31
- listings
 - add.c file, 200
 - add.s file, 200
 - annotation from chp10_profile 1 program, 273-276
 - annotation of one entry of cachegrind, 108
 - benefits profiling has on small programs, 15
 - call table for super block file system operations, 360
 - complex use of profiling, 26-27
 - creating a patch, 202
 - creating a thread with pthread_create, 72
 - demonstration of gdb's features, 57
 - disassembly of sys_open routine, 329
 - displaying commands available for crash using help command, 410
 - displaying crash_save_current_state routine using dis command, 403
 - displaying file information using files command, 404-405
 - displaying task information using task command, 405-409
 - dump file, 380-383
 - Electric Fence makefile for malloc_test.c, 94-95
 - Electric Fence malloc_test-fixed.c, 96
 - ext3 file system balloc.c source, 393-394
 - fat_clear_inode, 359
 - fix to bio.c, 205-207
 - fix to pnp_bios.c, 202-204
 - gcov after five tests, 50-51
 - gcov after running application with input 1000, 44-45
 - gcov output, 19, 25-26
 - gcov-enabled program, 42-43
 - gdbinit start file for gdb, 306, 343
 - getting a back trace using bt command, 396-397
 - gprof output, 17-21
 - high-level usage for adding a Kprobe, 318-319
 - ide_hwgroup_t structure, 366
 - ide_timer_expiry routine, 368
 - IRQ trace points, 235
 - jfs_clear_inode, 359
 - jfs_mount.c code, 344
 - Jprobe for sys_open routine, 330
 - Kprobe for sys_open routine, 325-326

- logrotate.conf file, 221-222
- makefile for sampleclock program, 6-7
- measuring cache misses with oprofile, 284
- MEMWATCH memory sample, 84-85
- modified jfs_mount.c code, 191-192, 390
- Oops directly from /var/log/messages, 195
- Oops messages after processing by ksymoops, 196
- oprofile sample program, 266-267
- portion of validate_super() routine, 105
- ps command showing a deadlock in JFS, 165-166
- ps-test.c, 157
- ps-test2.c, 160
- pseudocode for timing write code, 12-13
- sampleclock, 6-7, 9
- sampletime code, 11-12
- script to pull kernel and modules to test machine (kgdb), 340
- source for Syslog message, 224-226
- sqrt function, 20-21
- start of the strace on mkfs, 173
- Syslog message, 223-224
- syslog.conf file, 215, 217
- system call, 167
- sys_open source code, 323-324
- tar.data file, 254-255
- valgrind-1.c, 101
- valgrind-2.c, 103-104
- viewing a dump with crash, 394-395
- viewing file systems mounted using mount command, 399
- viewing system buffer information using log command, 400-403
- viewing system processes using ps command, 397-398
- viewing system run queue using runq command, 403
- viewing the back trace for process 1960 using bt command, 399
- viewing the machine information using mach command, 396
- viewing the system information using sys command, 395-396
- YAMD memory sample, 87
- LKCD (Linux Kernel Crash Dump), 372**
 - crash dump kernel options, 374-376
 - evaluating support for, 374
 - general crash dump steps, 379-384
 - lcrash commands, 385-386
 - patching and building the kernel, 376-378
 - sample system panic and crash dump taken, 386-390
- In commands, 190**
- Load all symbols for debugging option, 335**
- loadavg files, 129**
- log command, 400-403**
- log files, 215**
 - management (Syslog), 221
 - event logging, 228
- log records, 228**
- logger command, 218-219**
- logic errors, code coverage analysis, 39-40**
- logrotate command, 221**
- loop coverage, 40**
- loop grouping, performance counters, 280-281**
- lsdev files, 128**
- lsdf tool, 176**
- lspci utility, 118, 121**
- lsusb utility, 118**
- LTT (Linux Trace Toolkit), 234, 330**
 - analyzing tracing processes, 245
 - architecture component tasks, 235
 - building the kernel, 240-242
 - building user-level tools, 242
 - data acquisition, 243-244
 - data recording, 244
 - data reviewing text tools
 - traceanalyze filename, 257
 - tracedump filename, 258-259
 - tradedcore filename, 256
 - interpretation of data, 248-253
 - package and installation, 236-240
 - tarring a subdirectory sample, 253-256
 - Web resources, 259
- Lustre File System, 310**
- M**
- mach command, 396**
- magic key sequence, back traces, 174-176**
- Magic SysRq, 145**
- make bzImage command, 265, 320**
- make clean command, 265, 320**
- make command, 99**
- make install command**
 - installing YAMD, 88
 - Valgrind installation, 99
- make modules command, 265, 321**
- make modules_install command, 265, 321**
- make xconfig ARCH=um command, 299, 303**
- make xconfig command, 263, 320**
- malloc() error conditions, 46**
- malloc() function, 83**
- man gprof command, 15**
- man ps command, 154**
- management, event logs (evlogmgr command), 228**
- Management console option (UML configuration), 297**
- mcore crash dump, 393-410**
- md command, 366**

meminfo files, 129
memory corruption, 82
memory leaks, 82
memory management debugging, 82
 Electric Fence package, 93-96
 errors, 83-84
 libc (library calls), 82
 calloc() function, 83
 free() function, 83
 malloc() function, 83
 realloc() function, 83
 MEMWATCH, 84-86
 Valgrind, 97
 cache profiling, 106-109
 finding memory errors, 101-106
 installation, 98-101
 limitations, 109
 Web resources, 81, 109-111
 YAMD package, 86, 91-93
memory overruns, 97
MEMWATCH, 84, 86
messages, *syslogd* daemon, 217
mount command, 344, 399
mount root fs mnt -o loop command, 304
Mozilla, 116
mpstat (/proc file system), 134
Multicast transport, UML packet exchange, 301
mutex objects, 156
MySQL error code, *perror* utility, 208-209
N
net directory (/proc file system), 135-136
Netdump, 372, 391-392
netstat command, 180-182
network debugging tools, 178
 arp command, 179
 etherreal, 180
 ifconfig utility, 178
 netstat command, 180-182
 tcpdump program, 179
Networking support option (UML configuration), 297, 310
nm command, 189
null modem serial cable setup, 341-343
O
o option (Omit Event Types), *tracevisualizer* program, 252
objdump command, 356
Oops messages (panic messages), 188-190
 analysis, 190-195, 372
 applying gdb to display *jfs_mount*, 199

 generation of assembly code by gcc `-s` option, 199-200
 kernel mailing lists
 2.5.1-pre5, 204-207
 3c59x network driver, 200-204
 processing using *ksymoops*, 196-198
 Web resources, 209
opannotate command, 273
opannotate utility, 268
opcontrol -init command, 282
opcontrol utility, 267
open dialog box (kprof), 33
OpenOffice, 116
operations, gdb, 54-55
opgprof utility, 268
opreport utility, 268
oprofile kernel profiler, 263
 building the kernel, 264-265
 hardware counters, 277
 aligning structures, 279
 blocking technique, 281-287
 cache misses, 278
 loop grouping, 280-281
 packing arrays, 280
 padding structures, 279
 oprof_start utility, 268
 profile steps, 268-272
 reports, 276-277
 saving data, 277
 single executable profiles, 272-273
 utilities, 267-268
op_help command, 282
overruns (memory), 97
P
p option (PID tracing), *tracevisualizer* program, 253
packing arrays, performance counters, 280
padding structures, performance counters, 279
panic messages. *See* **Oops messages**
patch command, 202, 296
patching kernel, 376-278
 kdb, 348-350
 kgdb, 336-338
 UML (User-Mode Linux), 296
 configuration options, 297-299
 gcov support, 302
 gprof support, 303
 Kernel hacking menu, 301
 transports, 299-301
 UML Network Devices menu, 299
pcap transport, UML packet exchange, 301
Perens, Bruce, 93

- performance counters, 277**
 - aligning structures, 279
 - blocking technique, 281-287
 - cache misses, 278
 - loop grouping, 280-281
 - packing arrays, 280
 - padding structures, 279
- performance tuning**
 - profiling methods, 2
 - clock()* function, 6, 9-10
 - date* command, 4
 - gettimeofday()* function, 11-13
 - gprof* profiler, 13-20, 26-32
 - stopwatch* method, 3
 - time* command, 5-6
 - steps, 2
- perorr command, 188, 208-209**
- pg gcc option (gprof profiler), 15-16**
- pgrep -u root httpd (pgrep utility), 168**
- pgrep -u root,best (pgrep utility), 168**
- pgrep bash command, 169**
- pgrep utility, 168-169**
- PID 2, back trace, 362**
- POSIX 1003.25 standard, 226**
- post_handler function, 323**
- pre_handler function, 323**
- print array command, 46**
- print array[0] command, 49**
- print command, 60**
- Print Screen key, 145**
- printf() statements, 54**
- printk() messages, 227**
- Probe Manager (Dprobes), 316**
- Process Analysis view (tracevisualizer program), 250**
- process-specific subdirectories (/proc file system entries), 114**
 - /proc/vmstat* file, 146
 - /sys/kernel* file, 140-143
 - /sys/vm* file, 144-145
 - buddyinfo* file, 117
 - bus directory, 118
 - cmdline* file, 121
 - cpuinfo* file, 122
 - dev/fs* file, 138
 - diskstats* file, 122-126
 - free* file, 134
 - interrupts* file, 127
 - kallsyms* file, 128
 - kcore* file, 128
 - kmsg* file, 128
 - loadavg* file, 129
 - lsdev* file, 128
 - meminfo* file, 129
 - mpstat*, 134
 - net directory, 135-136
 - slabinfo* file, 136
 - stat* file, 137
 - statm*, 116
 - sys directory, 138
 - sysrq-trigger* file, 145
 - vmstat*, 131-133
- processes**
 - task_struct* data structure, 152-153
 - tools, 154
 - pgrep* utility, 168-169
 - ps* (process status) program, 154-157, 165-167
 - pstree*, 169
 - top* utility, 170
- processor entry (cpuinfo file, /proc file system), 122**
- processors, hardware counters, 277**
 - aligning structures, 279
 - blocking technique, 281-287
 - cache misses, 278
 - loop grouping, 280-281
 - packing arrays, 280
 - padding structures, 279
- profiling, 2, 262**
 - clock()* function, 6, 9-10
 - date* command, 4
 - gettimeofday()* function, 11-13
 - gprof* profiler, 13-14
 - b* option, 17
 - field descriptions*, 28-31
 - for loops/if statements*, 20, 26
 - fprofile-arcs* option, 18
 - fiest-coverage* option, 18
 - kprof*, 31-32
 - pg gcc* option, 15-16
 - instrumentation profilers, 262
 - sampling profilers
 - oprofile*, 263-287
 - Prospect*, 287-288
 - stopwatch* method, 3
 - time* command, 5-6
 - Web resources, 289
- Profiling support menu, 264**
- programs**
 - compilation, 59
 - identifying crashes, 79
- Prospect profiler, 287-288**
- provider interface (event logging), 227**
- ps (List All Processes) kdb command, 361**
- ps (process status) program, 154-157, 165**
 - displaying syscall currently executed, 166
 - ps au* option, 154-156
 - starting a new process, 167
- ps au option (ps program), 154-156**
- ps command, 116, 245, 397-398**

ps-test2, 162
 pstree utility, 169
 pthread_mutex_lock() function, 156
 pthread_mutex_unlock() function, 156

Q-R

quit command, 48-49, 61, 410

R (runnable) state of tasks, 156

r option (syslogd daemon), 213

Raw Trace view (tracevisualizer program), 250

Real-time Clock option (UML configuration), 298

real-time performance, 2

- clock() function, 6, 9-10
- date command, 4
- gettimeofday() function, 11-13
- gprof profiler, 13-14
 - b option*, 17
 - field descriptions*, 28-31
 - for loops/if statements*, 20, 26
 - fprofile-arcs option*, 18
 - ftest-coverage option*, 18
 - kprof*, 31-32
 - pg gcc option*, 15-16
- stopwatch method, 3
- time command, 5-6

realloc() function, memory management, 83

receive (RX) errors, ifconfig utility, 178

recording data, LTT (Linux Trace Toolkit), 244

registering Kprobes, 322-323

register_probe interface, 322-323

remote logging

- event logging, 229-230
- syslogd daemon, 214

reports

- generation, crash dumps, 373
- oprofile kernel profiler, 276-277

resources (Web resources)

- /proc file system, 149
- crash dumps, 411
- Dprobes, 331
- event logging, 232
- GNU Debugger, 53, 80
- kdb debugger, 369
- kgdb debugger, 369
- LTT (Linux Trace Toolkit), 259
- memory checkers, 81, 109-111
- Oops messages, 209
- profiling, 289
- system tools, 185

reviewing data, LTT text tools, 256

- traceanalyze filename, 257
- tracedcore filename, 256
- tracedump filename, 258-259

root file system (UML), 293-295

root images (UML), 304, 311

rpm (UML), 293

rpm command, 16

run 1 command, 46, 49

run command, 56, 60-61, 74

run-yamd command, 92

runq command, 403

RX (receive) errors, ifconfig utility, 178

S

S (sleeping) state of tasks, 156

s domainlist option (syslogd daemon), 213

sampling profilers, 263

- oprofile, 263
 - building the kernel*, 264-265
 - hardware counters*, 277-287
 - profile steps*, 268-272
 - reports*, 276-277
 - saving data*, 277
 - single executable profiles*, 272-273
 - utilities*, 267-268
- Prospect, 287-288
- Web resources, 289

sar utility, 126

segmentation fault - core dumped messages, 79

selector field (syslog.conf configuration file), 214-215

Separate Kernel Address Space support option (UML configuration), 297

separate loops, performance counters, 281

serial cable setup, 341-343

serial console port, 352-353

serial lines, UML (User-Mode Linux), 310

server requirements, Netdump, 391

set array=0 command, 46

set array[0]=0 command, 49

severity code (syslog.conf selector field), 214

Seward, Julian, Valgrind, 97

sharp profiles, 13

shutting down UML (User-Mode Linux), 312

single executable profiles, 272-273

slabinfo file (/proc file system), 136

SLIP transport, UML packet exchange, 301

SLiRP transport, UML packet exchange, 301

slog_fwd command, 231

sourcename files, 42

spinlock, 202

sqrt function, 20-21

SRAM (static RAM), 106

stack frames, 55

starting OpenOffice, 116

stat command, 387

stat file (/proc file system), 137

- static RAM (SRAM), 106
- statm entry (/proc file system), 116
- Steinberg, Udo A., 204
- step command, 48-49
- steps
 - crash dumps, 379-380, 384
 - profiling, 268, 272
- stopwatch, 3
- strace, 170-174
- strip command, 78
- stty command, 341
- Support for host-based filesystems option (UML configuration), 297
- Support for USB Keyboard in KDB option, 335
- SUSE Linux Enterprise Server, LKCD
 - functionality, 372
- swap partition, UML (User-Mode Linux), 306
- swap_fs swap partition, 306
- symbols (debugging), 77
- Symmetric multi-processing support option (UML configuration), 298
- sys command, 395-396
- sys directory (/proc file system), 138
- sysctl functions, 112
- Syslog, 212
 - changing messages, 217
 - debugging using messages, 222-226
 - disk space, 219-220
 - klogd, 217-218
 - log file management, 221
 - logger command, 218-219
 - options, 213
 - remote logging, 214
 - syslog.conf configuration file, 214-215
- syslog() messages, 227
- syslog.conf configuration file, 214-215
- syslogd daemon, 213
 - changing messages, 217
 - debugging using messages, 222-226
 - disk space, 219-220
 - klogd, 217-218
 - log file management, 221
 - logger command, 218-219
 - options, 213
 - remote logging, 214
 - syslog.conf configuration file, 214-215
- SysRq key, 145
- sysrq-trigger file (/proc file system), 145
- system call, 167
- system error messages
 - config files, 189
 - kernel files, 188-189
- Oops messages
 - analysis, 190-195
 - applying gdb to display jfs_mount, 199
 - generation of assembly code by gcc -s option, 199-200
 - kernel mailing lists, 200-207
 - processing using ksymoops, 196-198
 - Web resources, 209
- perror command, 208-209
- system.map, 189-190
- system information (/proc file system), 112
 - /proc/vmstat file, 146
 - /sys/kernel file, 140-143
 - /sys/vm file, 144-145
 - administrative applications, 113
 - buddyinfo file, 117
 - bus directory, 118
 - cmdline file, 121
 - cpuinfo file, 122
 - dev/fs file, 138
 - diskstats file, 122-126
 - entries, 114-116
 - free file, 134
 - interfaces, 112
 - interrupts file, 127
 - kallsyms file, 128
 - kcore file, 128
 - kmsg file, 128
 - loadavg file, 129
 - lsdev file, 128
 - meminfo file, 129
 - mpstat, 134
 - net directory, 135-136
 - relationship with sysctl functions, 112
 - slabinfo file, 136
 - stat file, 137
 - sys directory, 138
 - sysrq-trigger file, 145
 - vmstat, 131-133
 - Web resources, 149
- system panic, crash dump, 386-390
- system tools, 152
 - crash, 373
 - lcrash, 373
 - lsof, 176
 - magic key sequence, 174-176
 - network debugging tools, 178
 - arp command, 179
 - ethereal, 180
 - ifconfig utility, 178
 - netstat command, 180-182
 - tcpdump program, 179

- processes, 152-154
 - pgrep utility*, 168-169
 - ps (process status) program*, 154-157, 165-167
 - pstree*, 169
 - top utility*, 170
- strace, 170-174
- task states, 153
- Web resources, 185
- system.map, 189-190

T

- tap option (netstat command)**, 182
- tar command**, 98
- tarring a subdirectory sample (LTT)**, 253-256
- task command**, 405, 409
- task states**, 153
- TASK_DEAD**, 153
- TASK_INTERRUPTIBLE**, 153
- TASK_RUNNING**, 153
- TASK_STOPPED**, 153
- task_struct data structure**, 152-153
- TASK_UNINTERRUPTIBLE**, 153
- TASK_ZOMBIE**, 153
- TCP (Transmission Control Protocol)**, 230
- tcpdump program**, 179
- test machines**, 336
- test suites**, 38
- testing UML file systems**, 309
- thread 1 command**, 75
- thread 2 command**, 75
- thread 3 command**, 75
- thread 4 command**, 75
- thread analysis**, 308, 347
- thread apply command**, 309
- thread command**, 309
- time command**, 5-6
- TIMER_INT counter**, 281
- tools**, 152
 - crash, 373
 - lcrash, 373
 - lsof, 176
 - magic key sequence, 174-176
 - network debugging tools, 178
 - arp command*, 179
 - ethereal*, 180
 - ifconfig utility*, 178
 - netstat command*, 180-182
 - tcpdump program*, 179
 - processes, 152-154
 - pgrep utility*, 168-169
 - ps (process status) program*, 154-157, 165-167
 - pstree*, 169
 - top utility*, 170

- strace, 170-174
- task states, 153
- Web resources, 185
- top utility**, 170
- trace facility**, 234
 - applications, 234
 - LTT (Linux Trace Toolkit), 234
 - analyzing tracing processes*, 245
 - architecture component tasks*, 235
 - building the kernel*, 240-242
 - building user-level tools*, 242
 - data acquisition*, 243-244
 - data recording*, 244
 - data reviewing text tools*, 256-259
 - interpretation of data*, 248-253
 - package and installation*, 236-240
 - tarring a subdirectory sample*, 253-256
 - Web resources*, 259
- trace seconds filename script**, 243
- traceanalyze filename script**, 257
- tracecore seconds filename script**, 244
- tracecpuid seconds filename script**, 244
- tracedaemon command**, 242
- tracedcore filename script**, 256
- tracedump filename script**, 258-259
- traceu filename script**, 244
- tracevisualizer command**, 242, 251
- tracevisualizer program, LTT data interpretation**, 248-249
 - a option (Account Time Tracing), 253
 - c option (CPU ID tracing), 253
 - Event Graph, 249
 - f option (Event Types Tracing), 252
 - o option (Omit Event Types), 252
 - p option (PID tracing), 253
 - Process Analysis view, 250
 - Raw Trace view, 250
- Tracing thread support option (UML configuration)**, 297
- Transmission Control Protocol (TCP)**, 230
- transmit (TX) errors**, 178
- transports (UML packet exchange)**, 299-301
- TUN/TAP transport, UML packet exchange**, 300
- TX (transmit) errors**, 178

U

- UDP (User Datagram Protocol)**, 230
- ulimit command**, 63
- UML (User-Mode Linux)**, 292
 - applications, 292
 - booting, 307
 - consoles and serial lines, 310

- file systems, 305
 - adding*, 311
 - development and testing*, 309
 - gdb sessions, 306-308
 - gdb thread analysis, 308
 - introduction to kernel and root file system, 293-295
 - limitations, 293
 - networking support, 310
 - patching and building kernel, 296-304
 - configuration options*, 297-299
 - gcov support*, 302
 - gprof support*, 303
 - Kernel hacking menu*, 301
 - transports*, 299-301
 - UML Network Devices menu*, 299
 - root images, 304, 311
 - shutting down, 312
 - utilities, 312-313
 - Web resources, 314
 - UML Network Devices menu, 299
 - uml_mconsole utility, 312
 - uml_mkcow utility, 313
 - uml_moo utility, 313
 - uml_net utility, 313
 - uml_switch utility, 313
 - umount mnt command, 304
 - uname -a command, 265
 - uname command, 321
 - unregistering Kprobes, 322-324
 - unregister_probe interface, 323-324
 - User Datagram Protocol (UDP), 230
 - user-level tools, LTT (Linux Trace Toolkit), 242
 - User-Mode Linux. *See* UML
 - utilities
 - /proc file system, 114
 - /proc/vmstat file*, 146
 - /sys/kernel file*, 140-143
 - /sys/vm file*, 144-145
 - buddyinfo file*, 117
 - bus directory*, 118
 - cmdline file*, 121
 - cpuinfo file*, 122
 - devfs file*, 138
 - diskstats file*, 122-126
 - free file*, 134
 - interrupts file*, 127
 - kallsyms file*, 128
 - kcore file*, 128
 - kmsg file*, 128
 - loadavg file*, 129
 - lsdev file*, 128
 - meminfo file*, 129
 - mpstat*, 134
 - net directory*, 135-136
 - slabinfo file*, 136
 - stat file*, 137
 - statm*, 116
 - sys directory*, 138
 - sysrq-trigger file*, 145
 - vmstat*, 131-133
 - event logging, 228-229
 - oprofile kernel profiler, 267-268
 - UML (User-Mode Linux), 312-313
- ## V
- Valgrind, 97
 - cache profiling, 106-109
 - finding memory errors, 101-106
 - installation, 98-101
 - limitations, 109
 - valgrind ls --all install command, 100
 - validate_super() routine, 105
 - VFS (Virtual File System), 359
 - VM (virtual memory) systems, 292
 - vmlinux file, 188
 - vmlinux file, 188
 - vmstat (/proc file system), 131-133
- ## W-Z
- Web resources
 - /proc file system, 149
 - crash dumps, 411
 - Dprobes, 331
 - event logging, 232
 - GNU Debugger, 53, 80
 - kdb debugger, 369
 - kgdb debugger, 369
 - LTT (Linux Trace Toolkit), 259
 - memory checkers, 81, 109-111
 - Oops messages, 209
 - profiling, 289
 - system tools, 185
 - UML (User-Mode Linux), 314
 - where command, 75
 - x option (iostat utility), 124
 - YAMD package, 86-93