



INDEX

Symbols

`${HOME}/.k5login`, 286
`${HOME}/.rhosts`, 281
`${HOME}/.shosts`, 295
`${HOME}/.ssh`, 290, 292–294
`${HOME}/.ssh/authorized_keys`, 293
`${HOME}/.ssh/authorized_keys2`, 442
`${HOME}/.ssh/config`, 295
`${HOME}/.ssh/environment`, 295
`${HOME}/.ssh/identity`, 290
`${HOME}/.ssh/rc`, 295
`/boot`, 183, 192–193, 423, 496
`/boot/grub`, 182–183
`/boot/grub/device.map`, 183
`/boot/grub/grub.conf`, 182
`/dev`, 189, 195, 206
`/dev/console`, 314
`/etc`, 414
`/etc/cron.daily`, 321
`/etc/cron.daily/logrotate`, 321
`/etc/crontab`, 321
`/etc/dhclient.conf`, 307
`/etc/dhclient-enter-hooks`, 309
`/etc/dhclient-exit-hooks`, 309
`/etc/dhcpd.conf`, 307, 310
`/etc/exports`, 343, 405, 411
`/etc/flamethrower/flamethrower.conf`, 392
`/etc/fstab`, 190, 207, 209–210, 414–415
`/etc/group`, 258, 261, 268–269, 272
`/etc/gshadow`, 269
`/etc/hosts`, 220, 224, 229, 248–249, 263–264, 348, 480, 501
`/etc/hosts.allow`, 281, 334, 410
`/etc/hosts.deny`, 281, 334, 410
`/etc/hosts.equiv`, 281
`/etc/init.d`, 184, 243, 315
`/etc/init.d/functions`, 315
`/etc/init.d/krb5kdc`, 275
`/etc/init.d/nfs`, 405
`/etc/init.d/nfslock`, 405
`/etc/init.d/rcS`, 358, 382
`/etc/initlog.conf`, 315
`/etc/inittab`, 245, 346, 358
`/etc/krb5.conf`, 274, 323
`/etc/ld.so.conf`, 464
`/etc/log.d`, 321
`/etc/log.d/logwatch.conf`, 321
`/etc/logrotate.conf`, 319–320



/etc/logrotate.d, 320
/etc/logrotate.d/kadmind, 323
/etc/logrotate.d/krb5kdc, 323
/etc/logrotate.d/named, 320
/etc/logrotate.d/syslog, 320
/etc/mdadm.conf, 197, 202–203
/etc/modules.conf, 238–239
/etc/named.conf, 250
/etc/nologin, 295
/etc/nsswitch.conf, 224, 230, 248, 256, 259, 263, 265, 414
/etc/ntp.conf, 245, 308
/etc/pam.conf, 269
/etc/pam.d, 269
/etc/passwd, 258, 261, 263, 268–269, 272, 295
/etc/raidtab, 188, 190, 194–195, 197, 200, 202
/etc/rc.sysinit, 245
/etc/resolv.conf, 224, 230, 254, 308
/etc/rndc.conf, 252
/etc/services, 317, 461
/etc/shadow, 269
/etc/smartd.conf, 215
/etc/ssh, 287, 289
/etc/ssh/hosts.equiv, 295
/etc/ssh/known_hosts, 292
/etc/ssh/ssh_config, 290, 295
/etc/ssh/sshd_config, 289, 295–296, 348
/etc/ssh/sshrcc, 295
/etc/sshd_config, 348
/etc/sysconfig, 237, 257, 265
/etc/sysconfig/hwconf, 366
/etc/sysconfig/network, 259, 498
/etc/sysconfig/network-scripts, 237, 240, 242, 308
/etc/sysconfig/nfs, 407–408, 410
/etc/sysconfig/ntpd, 245
/etc/sysconfig/syslog, 316–318
/etc/sysctl.conf, 409–410
/etc/syslog.conf, 317
/etc/systemimager, 356
/etc/systemimager/flamethrow-
er.conf, 392, 394
/etc/systemimager/flamethrow-
erd.conf, 393
/etc/systemimager/image-
man.perm, 394
/etc/systemimager/imagema-
nip.conf, 394
/etc/systemimager/px-
elinux.cfg, 357
/etc/systemimager/rsync_stubs/
99local, 380, 385
/etc/systemimager/rsyncd.conf,
380, 385
/etc/xinetd.conf, 283
/etc/xinetd.d, 283
/etc/xinetd.d/eklogin, 283
/etc/xinetd.d/gexec, 461
/etc/xinetd.d/tftp, 334
/etc/yp.conf, 260, 308
/export, 496
/export/home/install, 497
/home, 489
/images/pxeboot, 343
/lib, 185
/lib/modules/2.4.25-udpcast, 373
/lib/security, 269
/proc, 205, 384, 410
/proc/cmdline, 380, 384
/proc/kmsg, 316
/proc/mdstat, 189, 193, 199
/proc/scsi/scsi, 201
/proc/sys/dev/raid/
speed_limit_max, 197
/proc/sys/dev/raid/
speed_limit_min, 190, 197
/proc/sys/net/core/
rmem_default, 410
/proc/sys/net/core/rmem_max, 410
/root/.ssh, 375
/root/anaconda-ks.cfg, 340
/sbin/dhclient, 307
/sbin/dhclient-script, 308
/sbin/mkbootdisk, 177
/sbin/mkinitrd, 185

- /tftpboot, 339, 343, 357
 - /tftpboot/pxelinux.0, 312
 - /tftpboot/pxelinux.cfg, 337
 - /tftpboot/pxelinux.cfg/default.netboot, 378, 382, 385
 - /tmp, 291
 - /usr/bin, 289
 - /usr/bin/ssh-keygen, 288
 - /usr/include/sys/syslog.h, 318
 - /usr/kerberos/bin, 279, 282
 - /usr/kerberos/sbin, 282
 - /usr/lib/mpich-1.2.5.2/sbin/chp4_servs, 443
 - /usr/lib/systemimager, 356
 - /usr/lib/udpcast/templ, 373
 - /usr/lib/yp, 256
 - /usr/local/lib, 464
 - /usr/local/nagios, 464
 - /usr/local/nagios/etc, 466
 - /usr/local/nagios/libexec, 465
 - /usr/local/nagios/share/images/logos, 472
 - /usr/sbin, 423
 - /usr/sbin/dhcpd, 307
 - /usr/sbin/logwatch, 321
 - /usr/sbin/rpc.mountd, 405
 - /usr/sbin/rpc.nfsd, 405
 - /usr/sbin/sshd, 288–289
 - /usr/sbin/udp-receiver, 373
 - /usr/sbin/udp-sender, 373
 - /usr/share/boot/systemimager, 380
 - /usr/share/hwdata, 366
 - /usr/share/systemimager/i386/standard, 390
 - /var, 192
 - /var partition, 319
 - /var/kerberos/krb5kdc, 275
 - /var/kerberos/krb5kdc/kadm5.acl, 276
 - /var/kerberos/krb5kdc/kadm5.keytab, 276
 - /var/kerberos/krb5kdc/kdc.conf, 274–275
 - /var/kerberos/krb5kdc/kprop.acl, 278
 - /var/kerberos/krb5kdc/kpropd.acl, 278
 - /var/lib/dhcp, 307
 - /var/lib/dhcp/dhclient-eth0.leases, 307
 - /var/lib/logrotate.status, 321
 - /var/lib/mysql/oscar, 489
 - /var/lib/nfs/rmtab, 411
 - /var/lib/nfs/xtab, 411
 - /var/lib/systemimager, 356
 - /var/lib/systemimager/images, 357
 - /var/lib/systemimager/scripts, 359, 377, 382, 390, 486
 - /var/log/boot.log, 315
 - /var/log/kadmind.log, 280
 - /var/log/krb5kdc.log, 276, 280, 284
 - /var/log/messages, 178, 216, 307, 323, 337
 - /var/log/messages/flamethrower.flamethrower_directory, 392
 - /var/log/systemimager/rsyncd, 385
 - /var/named, 252
 - /var/yp, 257
 - /var/yp/Makefile, 260–262
- A**
- Accelerated graphics port, 91
 - ACL
 - See Authorization, access control list
 - Advanced Research Projects Agency, 282–283
 - AFS
 - See File system, Andrew
 - AGP
 - See Accelerated graphics port
 - API
 - See Application programming interface
 - Application
 - parallel
 - NWChem, 65

Application programming interface, 21–22, 62
 Globus toolkit, 506
 load balancing, 438
 Maui scheduler, 449
 Myrinet GM, 148
 OpenMOSIX, 504
 remote procedure call, 13
 socket, 149

ARPA
See Advanced Research Projects Agency

Authentication
 application execution, 434
 encrypted transport, 287
 for Maui jobs, 448
 Ganglia authorization daemon, 451
 group identifier, 212, 261, 268, 270, 348,
 411, 448, 464, 504
 Kerberos, 272, 325
 for BSD/ARPA services, 282
 mutual, 287
 network information service, 263
 Network Time Protocol, 245
 password, 269
 pluggable authentication modules, 269,
 279, 446
 public key, 288
 remote access, 281
 secure shell, 229, 286
 authorized keys file, 293
 shadow files, 261
 user, 51
 user identifier, 261, 268, 270, 348, 411,
 448, 464, 504
 user identifier and group identifier, 268
 X-Windows, 297

Authorization
 access control lists, 268
 directories
 default permissions, 270
 set group identifier bit, 270–271
 sticky bit, 270
 umask, 271

autofs
See File system, automounter

B

Basic input output subsystem, 36, 71, 120,
 175, 177–178, 181–182, 198
 flashing of, 175
 history of, 175
 interface card, 176
 partitioning, 176, 198
 upgrading, 72, 535, 539–540

Benchmark
 CFP2000, 79
 CINT2000, 79
 Linpack, 74
 SPEC CPU2000, 79
 SPECfp_2000, 79
 SPECfp_rate_2000, 79
 SPECint_2000, 79
 SPECint_rate_2000, 79

BIOS
See Basic input output subsystem

Booting
 initial RAM disk image, 184, 186, 188,
 191, 194–195, 197, 205, 334,
 339, 343, 346, 357, 367, 373,
 377–378, 380–382, 389, 393
 isolinux, 334
 pre-execution environment, 307–308,
 312, 333, 367
 pxelinux, 220, 230, 334, 337, 339, 343,
 357, 366, 377, 379, 381, 384,
 389
 activating, 339
 configuration, 339
 configuring, 336
 syslinux, 334

Bootloader
 elilo.efi, 181–182
 extensible firmware interface, 182
 grub, 182–184, 334
 lilo, 182, 334

BOOTP
See Network, protocol, Internet Bootstrap
 Protocol

British thermal unit, 34, 557, 563–564

BTU

See British thermal unit

C

Cable

classes of, 546
connection point identifier, 550, 556
 example of, 549, 551
fiber-optic, 54, 96, 98, 110, 139, 146,
 149, 153, 546, 548
Infiniband reduction of, 514
interrack, 28, 36–37, 49, 142, 519, 537–
 538, 541–542, 548, 551–553
 definition of, 547
 minimizing, 49, 531, 553
intrarack, 37, 541–542, 548, 551–553
 definition of, 547
management of, 545
number of, 513
serialized, 551
shielded twisted pair, 96, 110

Cable management

rack-mounted arms, 36

Cache memory

latency in, 87

Cache-coherent nonuniform memory access,
22, 47, 76

ccNUMA

See Cache-coherent nonuniform memory
access

CDMA/CD

See Ethernet, carrier detect multiple ac-
cess with collision detect

Checkpoint and restart, 61, 129, 398, 404, 504

CIFS

See File system, common Internet

Cluster

assembly
 necessary tools, 530
 rules of thumb, 530
 step-by-step process, 532
categories of, 23

complexity, 521

compute rack, 546

creation of, 28

definition of, 22, 24

design of, 30

hardware elements, 45

 active, 47–48, 525

 passive, 46

heat

 cold aisles, 562

 conversion factors, 563

operational testing, 39

performance scaling across, 81

physical assembly of hardware, 36

power

 conversion factors, 561

 manufacturer's specifications, 560

 power factor, 562

 total requirements, 560

preliminary definition of, 4

resource management system, 437

 Quadrics, 437

 simple Linux utility for, 445

type

 carpet, 15, 52

 high availability, 23, 68

 high performance, 23, 68

 high throughput, 23, 68

Clustered system

See Cluster

Community

open-source, 161

Complex instruction set computing, 9

Compute

farm, 53

ranch, 53

Compute slice, 25, 30, 32, 37, 46–48, 58, 69,
80, 119, 546

1U versus 2U, 512

64-bit capabilities, 82

64-bit software, 85

cabling, 546

cloning, 330

data access, 397

- database instances, 434
 - definition of, 47
 - diskless, 345
 - economics, 511
 - example clusters, 80
 - failure and replacement, 545
 - graphics, 91
 - hardware architecture, 73
 - HSI performance, 134
 - I/O bus, 90
 - I/O capacity, 89
 - image capture, 231
 - installation, 228, 267, 327, 352, 366
 - local storage, 217
 - memory bandwidth, 72, 86
 - memory hierarchy, 87
 - memory latency, 87
 - number of processors, 89
 - operating system choice, 165
 - peak floating point, 74
 - physical RAM capacity, 83
 - power utilization, 560
 - relative cost, 79
 - resource management, 437
 - resource monitoring, 455
 - scaling up with, 160
 - selection criteria, 70
 - total system cost, 78
 - user access, 446
- Connection point ID
See Cable, connection point identifier
- Console
video graphics array, 121
- CPAN
See Software, comprehensive Perl archive network
- Custer
- context diagram, 30
 - heat
 - hot aisle, 562
 - refrigeration ton, 563
 - installation process, 30
 - management rack, 546
 - physical assembly of, 529
 - site preparation, 34
 - testing
 - process of, 30
- ## D
- Data
- computational independence of, 14
 - Ganglia distributed monitoring, 449
- Data transmission
- duplex
 - debugging, 242
 - full duplex, 96, 110, 135, 140
 - definition of, 109
 - half duplex, 96, 110
 - definition of, 109
- Database
- application demands, 398
 - authentication, 269
 - cluster, 61, 149
 - cluster configuration, 249
 - direct access mode, 428
 - domain name service, 252, 255
 - Ganglia metric, 455, 457
 - grid access, 506
 - hardware identifiers, 366
 - host name lookup, 249
 - instance, 48, 434
 - Kerberos, 275
 - Kerberos key distribution center, 276, 278–279
 - MySQL, 446, 478, 489, 502, 578
 - name service switch selection of, 249
 - network information service, 248, 256–257, 259, 261–262
 - open-source, 578
 - Oracle software, 431
 - parallel, 30, 51, 59–60, 134, 434
 - resource management, 438
 - round robin, 450, 452–453, 578
 - secure shell known hosts, 293
 - server, 15, 50, 57, 89, 125, 133, 158, 280, 432, 435

- storage, 130
- system management, 479
- `SystemConfigurator` hardware mapping, 366
- tables, 398
- transaction processing, 16
- `tripwire`, 227
- DCE
 - See* Distributed computing environment
- Debugging
 - core dump, 268
 - crash cart, 121
 - devices with `/proc`, 205
 - Kerberos key distribution center, 283
 - Nagios, 473
 - network information service, 263
 - network issues, 37, 115
 - parallel debugger, 474
 - software issues, 313
 - system logging for, 313
 - `SystemImager` image capture, 362
- Dependencies
 - loosely coupled, 16
 - tightly coupled, 17
- Device
 - driver, 38, 134, 206, 214, 444
 - file, 178–179, 189, 195, 206, 209, 213, 356
 - `mknod`, 206
 - Firewire, 180, 201, 206, 421
- Device management
 - out of band, 119
- DHCP
 - See* Network, protocol, Dynamic Host Configuration Protocol
- Digital signal processors, 62
- Digital visual interface, 63
- Digital-to-analog converter, 63
- Direct memory access, 89
- Disk
 - default partitioning, 176
 - floppy format, 177
 - logical units, 195
 - logical volume manager, 217
 - master boot record, 177–178, 181–182
 - multiple device facility, 188
- Distributed computing environment, 11, 13, 272
- Divide et impera*, 28
- DNS
 - See* Domain name service
- Domain name service, 264
 - configuration file, 251
 - configuring `rndc`, 252
 - debugging with `rndc`, 255
 - forwarding, 252
 - record types, 252
 - reverse mapping, 252
 - testing, 253
 - zone definitions, 252
- E
- EDA
 - See* Electronic design automation
- EFI
 - See* Extensible firmware interface
- EIA
 - See* Electronics Institute of America
- Electronic design automation, 54
- Electronics Institute of America, 49, 70, 541
 - EIA unit, 81, 541, 546, 548
- Enabling services
 - `chkconfig`, 243
- EPIC
 - See* Explicitly parallel instruction set computing
- Error correction code, 72
- Ethernet, 9, 22, 34, 48, 52–53, 95–96, 160, 403, 478, 513
 - 10-gigabit, 96, 112, 117, 141–142, 145
 - Aloha Net, 96
 - autonegotiation, 110, 114, 140, 233, 242
 - bandwidth, 111, 116, 128, 140–143, 145, 237, 404
 - carrier detect multiple access with collision detect, 96

- frame, 98, 100–101, 106, 113, 241
 - gigabit, 5, 92, 96, 110, 117, 128, 134, 140, 142–143, 238, 363, 366, 370, 404, 429, 523, 535, 546, 550–551
 - hub, 99
 - IEEE standard, 96, 98, 112–113
 - jumbo frames, 98, 113, 128–129, 142, 241
 - maximum transmission unit, 113
 - media access control address, 98
 - network design with, 125
 - organizationally unique identifier, 98
 - router, 100
 - segment, 97, 100, 104, 107
 - switch, 109, 113, 115–116, 120, 122, 128, 131, 135, 139–140, 154, 230, 234, 238, 512, 521, 531, 533–534
 - switch meshing, 141
 - topology, 97
 - Explicitly parallel instruction set computing, 9
 - Extensible firmware interface, 181
 - applications, 181
 - boot loader, 181
 - elilo.efi, 182
 - system partition, 181, 193
 - Extensible mark-up language, 342, 360, 423, 425–427, 447, 450
 - External data representation, 450
- F**
- Fibre Channel, 129, 195, 404, 419, 428, 535
 - File system, 574
 - /proc, 342, 450
 - Andrew, 272
 - automounter, 231, 414–415, 502
 - chroot command, 334
 - cluster, 38, 58, 62, 123, 134, 204, 397–398
 - definition of, 204
 - IBM general parallel file system, 58
 - InterMezzo, 58
 - Lustre, 58, 204
 - Oracle, 204, 428, 431
 - common Internet, 10, 204
 - compressed image, 184
 - consistency, 59
 - context-dependent symbolic links, 431
 - data blocks, 223
 - DOMAIN/OS, 10
 - enterprise volume management system, 420
 - ext3, 188
 - creating, 209
 - journal, 207–208, 211
 - stride, 211
 - superblock, 210
 - volume label, 209–210
 - file allocation table, 177, 181, 188
 - highly available, 58
 - journal, 208
 - kernel modules, 185
 - local, 175, 188, 217, 399, 401
 - ext, 207
 - ext2, 205, 207
 - ext3, 205
 - iso9660, 205
 - Reiser, 205, 214
 - tempfs, 213, 384
 - Lustre, 421
 - example, 425
 - meta-data server, 421–423, 426
 - object storage server, 421, 423, 427
 - object storage target, 426
 - storage, 422
 - memory based, 346
 - meta-data, 59, 211, 370, 399–401
 - inode, 210
 - server, 59, 418
 - mounting options, 212
 - network, 58
 - remote procedure calls, 11
 - network attached storage, 400
 - networked, 10, 332, 404
 - definition of, 204
 - distributed file system, 398

- new technology file system, 204
 - Open Global File System, 419
 - OpenMOSIX, 503
 - parallel, 398, 401–403, 421
 - global file system, 428
 - Lustre, 134, 421
 - open-source, 415
 - PolyServe Matrix, 430–431
 - parallel virtual file system, 416, 418–419, 500, 574
 - meta-data server, 416
 - physical, 175, 204
 - definition of, 204
 - ext3, 207
 - PolyServe Matrix
 - meta-data, 430
 - POSIX semantics, 428
 - pseudo, 205–207
 - /proc, 205
 - /proc, 200
 - device file system, 206
 - remote root, 332
 - superblock, 184, 208
 - tuning, 214
 - virtual file system layer, 416
- Folding@home, 52
- FORTRAN, xxxvii, 441, 474
- Free Software Foundation, 161
- Front-side bus, 72
- ## G
- GARP
 - See* Network, protocol, Generic Attribute Registration Protocol
- GbE
 - See* Ethernet, gigabit
- Genome@home, 52
- GID
 - See* Authentication group identifier
- GMRP
 - See* Network, protocol, GARP Multicast Registration Protocol
- GNU
 - definition of, 161
 - GNU/Linux, 161
 - public license, 161
- GPFS
 - See* File system, cluster, IBM general parallel file system
- Graphical user interface, 15
- Graphics interchange format, 455
- Grid
 - definition of, 504
 - Globus Toolkit, 505
 - monitoring with Ganglia, 505
- ## H
- Hardware
 - power-on self-test, 121, 183, 540
- High-availability, 475
 - heartbeat, 57
- High-speed interconnect, 22, 25, 28, 30, 33, 37, 45–46, 48–49, 53, 63, 80, 89, 115, 122–123, 125, 127, 134–135, 137–139, 143, 145–146, 148, 154, 160, 231, 247, 333, 369, 416, 433, 435, 437, 444, 479, 492, 512–515, 517, 519, 521, 523, 531, 533–534, 538, 541–542, 546, 554, 576
 - Dolphin, 151, 154, 576
 - Ethernet, 122, 127, 133, 136, 139–140, 143, 145, 444, 446, 497, 512
 - gigabit Ethernet, 60
 - Infiniband, 22, 60, 148–151, 154, 576
 - Myrinet, 22, 48, 60, 138, 146, 148, 154, 416, 421, 440, 444, 446, 478, 492, 500–501, 507, 512, 519, 524, 534, 536, 546, 551, 576
 - dual-rail, 146
- Quadrics, xxxvii, 22, 66, 80, 152–154, 306, 416, 421, 437, 512, 516, 519, 524, 534, 576
- Quadrics QsNet, xxxvii
- topology, 137

- Host name
 - lexical ordering of, 247
- HSI
 - See* High-speed interconnect
- Hypertext
 - markup language, xxxviii, 279, 342, 352, 356, 441
 - preprocessor, 451, 456
- I**
- IEEE-1394
 - See* Device, Firewire
- IGMP
 - See* Network, protocol, Internet Group Management Protocol
- iLO
 - See* Management, integrated lights out
- init, 184, 186–187, 243, 315, 346, 380, 382, 384, 409
- initrd
 - See* Booting, initial RAM disk
- Installation
 - anaconda, 183, 191, 340, 342
 - Dynamic Host Configuration Protocol, 333
 - kickstart, 37, 199, 220, 332, 334, 339–341, 343, 349, 478
 - configuration file, 344
 - configuration kernel, 345
 - installation kernel, 343
 - kudzu, 366
 - rsync, 357
 - system personality, 331
 - system structure, 332
 - SystemConfigurator, 351, 365–366, 389
 - SystemImager, 37, 199, 220, 231, 332, 334, 339, 349, 351, 373, 375, 394, 479
 - flamethrower, 391, 393
 - server scaling, 370
 - updateclient, 389
 - SystemInstaller, 351
 - TFTP, 334
 - udpcast, 373
- Institute of Electrical and Electronics Engineers (IEEE), 96
 - organizationally unique identifier, 98
 - standards
 - 1394, 180
 - 1596-1992, 151
 - 802, 98
 - 802.1Q, 113
 - 802.3, 96
 - 802.3AD, 112
- IP
 - See* Network, protocol, Internet Protocol
- IPoIB
 - See* Network, protocol, IP over Infiniband
- IPv4
 - See* Network, protocol, Internet Protocol version 4, 107
- IPv6
 - See* Network, protocol, Internet Protocol version 6
- J**
- Java
 - application serving, 431
 - Cryptix security software for, 446
 - Globus grid toolkit, 506
 - inclusion in Red Hat Linux, 166
 - inclusion in Rocks, 493
 - Linux software, 576
 - management interface, 121
 - Maui scheduler's use of, 446
 - virtual machine, 446, 493
- Job execution
 - gexec, 450–451, 455, 460–462
 - mpirun, 435, 437, 441–443, 449
 - secure start shell, 448
- JVM
 - See* Java, virtual machine

K**KDC**

See Kerberos, key distribution center

Kerberos

administrative daemon, 276

database, 275, 277

kadmin

access control list, 276

kadmin server, 276

Kerberized services, 279

key distribution center, 230, 272

master, 278

slave, 278

name service requirements, 247

pluggable authentication modules, 279

principal database, 272

principals, 273

realm, 272

ticket granting service, 286

ticket granting ticket, 272, 279

time synchronization requirements, 244,
274

version 5, 272

version 5 on Linux, 274

Kernel direct-access programming library,
150

Kernel parameters

tuning, 410

tuning or setting, 409

L**LAM**

See Message-passing interface, local area
multicomputer

LAN

See Network, local area

LDAP

See Network, protocol, Lightweight
Directory Access Protocol

Linux Terminal Server Project, 346, 348, 504,
573

Linux virtual server, 57, 576

Load balancing, 54, 231, 435–436, 504

Linux virtual server, 436

open portable batch system, 436, 438,
441, 481, 489, 500, 502, 580

Platform Computing LSF, 54

portable batch system, 54

portable batch system Pro, 436

Sun Grid Engine, 438, 493

Lustre

See File system, Lustre

LVS

See Linux virtual server

M**Management**

integrated lights out, 71, 120

Mandelbrot set, 14

Marketecture, 53

Master node, 49, 92

redundancy, 92

MBR

See Disk, master boot record

MDS

See File system, Lustre, meta-data server

Mellanox Technologies Inc., 150

Memory

double data rate, 72, 86

dual in-line memory modules, 86

latency of, 87

single in-line memory modules, 86

synchronous dynamic random access, 86

synchronous dynamic random access
memory, 72

Message-passing interface, xxxvii, 21, 154,
160, 416, 435, 439–442, 446, 474, 481, 483,
492, 506, 579

chameleon, 148, 439–441, 443–446, 475,
481, 483, 502, 579

chameleon grid, 506

forum, 440

- local area multicomputer, 444, 446, 475, 481, 483, 489
 - standard
 - version 1.1, 440
 - version 2.0, 440
- Moore's Law, 9
- MPI
 - See* Message-passing interface
- MPICH
 - See* Message-passing interface, chameleon
- MTU
 - See* Network, maximum transmission unit
- MZAP
 - See* Network, protocol, Multicast Scope Zone Announcement Protocol
- N**
- Name resolution
 - host file, 249
- Name service switch, 224, 248
- NAS
 - See* Network attached storage
- NAT
 - See* Network, address translation
- National Partnership for Advanced Computational Infrastructure, 477
 - Rocks, 478–480, 492, 502
 - CD-ROMs, 493
 - front-end, 498
 - gathering MAC addresses, 500
 - general cluster information, 496
 - installation, 495
 - ISO images, 493
 - kickstart, 493
 - network configuration, 497
 - subsystem use, 501
 - system administration, 501
- Net mask
 - See* Network, subnet mask, 104
- Network
 - address classes, 104
 - address space, 104
 - address translation, 108
 - backbone, 101
 - bisection width, 137
 - bonding
 - See also* Network, trunked link, 238
 - booting
 - pxelinux, 334
 - class D address, 372
 - classless interdomain routing, 105
 - collision domain, 98
 - cross-section width, 137
 - cross-sectional bandwidth, 137–138
 - data LAN, 514
 - diameter, 136, 139
 - Dynamic Host Configuration Protocol
 - client, 237, 241, 308, 393
 - forward error correction, 375
 - frame
 - broadcast, 107
 - multicast, 107
 - unicast, 106
 - ifconfig, 236
 - interface
 - promiscuous mode, 97
 - Internet service provider, 105
 - ip command, 235
 - latency, 142
 - link-level control, 98
 - Linux bonding, 238
 - local area, 48–49, 61, 70, 96–99, 101, 106, 109–110, 113, 122, 127–128, 131, 140, 150–151, 223, 237, 247, 333, 371–372, 419, 513–514, 523, 531, 533–535, 538, 542, 546, 554–555
 - management LAN, 514
 - maximum transmission unit, 98, 241
 - setting the size of, 242
 - media access control, 98
 - media independent interface, 242
 - mii-tool, 243
 - multicast, 103, 371
 - administrative scoping, 372

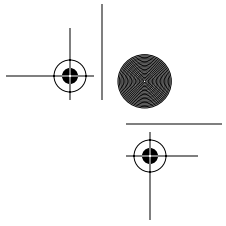
- group, 372
- group identifier, 372
- time to live, 371
- udpcast, 375–377, 386, 391, 394
- udp-receiver, 374, 379, 386, 388
- udp-sender, 374, 376, 386, 388
- nonblocking, 134
- open system interconnect, 96
- protocol
 - Address Resolution Protocol, 107, 236
 - Dynamic Host Configuration Protocol, 37, 108, 115, 120, 122, 159, 220, 223, 230, 232, 241, 260, 265, 307–309, 311, 325, 331, 333, 336–337, 339–340, 342, 346, 357, 361, 367, 369, 479, 488, 498, 501
 - GARP Multicast Registration Protocol, 371
 - Generic Attribute Registration Protocol, 371
 - Internet Bootstrap Protocol, 307, 333, 339, 367
 - Internet Group Management Protocol, 371
 - Internet Protocol, 102
 - Internet Protocol version 4, 107
 - Internet Protocol version 6, 108
 - IP over Infiniband, 150
 - Lightweight Directory Access Protocol, 269
 - Multicast Scope Zone Announcement Protocol, 372
 - Network Time Protocol, 232
 - Network Time Protocol, 122, 220, 229–231, 244–245, 262
 - remote direct memory access, 96, 150
 - remote monitoring, 115
 - SCSI remote DMA protocol, 150
 - server message block, 10, 204, 269
 - Simple Network Time Protocol, 115
 - Socket Direct Protocol, 149–150
 - Spanning Tree Protocol, 140
 - Transmission Control Protocol, 102
 - Transmission Control Protocol/Internet Protocol, 96, 98, 102, 116, 134, 148–149, 160, 236, 287, 296, 371, 403, 405, 416, 443–444, 476, 487, 505
 - Trivial File Transfer Protocol, 122, 220, 230, 232, 334, 336–337, 339, 343, 355
 - client, 336
 - User Datagram Protocol, 102
- rearrangeable, 138
- secure socket layer, 115
- subnet mask, 104
- supernet, 106
- supernetting, 123
- switching domain, 110
- topology, 133
 - bus, 136
 - Clos, 138, 146
 - completely connected, 136
 - fat tree, 139–140
 - flat network neighborhood, 145
 - hypercube, 136
 - mesh, 136
 - star, 136
 - 3D hypercube, 136
 - toroidal, 136
- Transmission Control Protocol wrapper, 281
- trunked link, 143, 238, 554
 - configuration of, 238
- virtual local area, 113–114, 116–117, 120, 125, 127–129, 144, 233–234
- wide area, 100–102, 414, 505
- network
 - Dynamic Host Configuration Protocol client, 307
- Network attached storage, 59

- Network file system, 10, 123, 134, 204, 263, 332, 342, 346, 401, 404, 415
 - automounter, 414
 - client subsystem, 412
 - cluster NFS, 348
 - diskless, 346
 - exports, 405, 411
 - mount parameters, 412, 414
 - mounting, 410
 - server, 92
 - server subsystem, 412
- Network information service, 256, 263
 - automounter maps, 225, 231, 248, 256, 264, 414
 - client, 260
 - domains, 257
 - maps, 256
 - master server, 256
 - slave server, 256
 - slave servers, 259
 - testing, 262
- NFS
 - See* Network file system
- NIS
 - See* Network information service
- NPACI
 - See* National Partnership for Advanced Computing Infrastructure
- NPACI Rocks
 - See* National Partnership for Advanced Computational Infrastructure, Rocks
- NTP
 - See* Network, protocol, Network Time Protocol
- O**
- OCFS
 - See* File system, cluster, Oracle
- Open multiprocessing, 21, 154, 416, 474
- OpenMOSIX, 502–503
 - booting from CD-ROM, 504
 - file system, 503
 - programming on, 504
 - unique home node, 503
- OpenMP
 - See* Open multiprocessing
- OpenPBS
 - See* Load balancing, open portable batch system
- Open-source cluster application resources, 352, 477, 479–480, 507, 573
 - as a general process, 491
 - cluster flavors, 480
 - configuration of, 480
 - features, 478
 - installation
 - client, 485
 - completing, 489
 - Dynamic Host Configuration Protocol, 488
 - MAC address collection, 488
 - SystemImager, 486
 - wizard, 481, 485
 - message-passing, 489
 - package, 481–482
 - configuration, 486
 - downloading, 482
 - server, 485
- Open-source software
 - See* Software, open-source, 170
- Oracle
 - 9i Real Application Cluster (RAC), xxx-viii, 22, 59, 428, 431, 476
 - database clients, 431
 - remote procedure call, 431
 - software qualification, 165, 173
 - using PolyServe Matrix, 430
- OSCAR
 - See* Open-source cluster application resources
- OSS
 - See* File system, Lustre, object storage server
- OSSD
 - See* Software, open-source, development

- OST
 See File system, Lustre, object storage target
- OUI
 See Ethernet, organizationally unique identifier
- P**
- Pacific Northwest National Laboratory (PNNL), xxxv, xxxvii–xxxviii, 23, 34, 65, 152, 373
- Parallel application
 debugging, 13, 474
 simplistic example of, 18
- Parallel distributed shell, 220, 229, 293, 300–304, 306, 367, 390
- Parallel virtual machine, 21
- PDSH
 See Parallel distributed shell
- PDU
 See Power distribution unit
- Peripheral component interconnect, 71, 89
 extended, 71
 implementation, 90
- PID
 See Process identifier
- PNNL
 See Pacific Northwest National Laboratory, xxxv
- Portable network graphics, 455
- Power distribution unit, 546
- Power supply
 rating caveats, 72
- Price-to-performance ratio, 79
- Process
 as an insult, 29
 repeatable cluster creation, 29
- Process identifier, 186, 315, 392
- Processor
 AMD Opteron 246, 78
 floating point, 6
 Intel Itanium 2 6M, 78
 Intel Pentium 4 Xeon, 78
- Processor architecture
 AMD x86_64, 478
 AMD32, 168–169
 AMD64, 167–169, 171
 explicitly parallel instruction set computing, 80
 Intel IA-32, xxxvii, 70, 83, 158, 168, 171–172, 182, 210, 521, 528
 Intel IA-64, xxxvii, 172, 181, 478
 Intel x86, xxxvii, 169, 171–172, 478
- Processor family
 AMD Opteron, 76–77, 79–82, 85, 89, 91, 158, 167–168, 478, 575
 HP Alpha, 158, 172
 HP PA-RISC, 158, 172, 528
 Intel Itanium, xxxvii–xxxviii, 9, 34, 66, 75, 77, 79, 81–82, 85, 91, 154, 158, 167–168, 172, 181–182, 193, 351, 355, 366, 373, 478, 491, 507, 512, 517, 521, 527–528, 573–574
 Intel Pentium, xxxvii, 70, 72, 74, 78, 81–82, 85, 512, 517, 561
 Intel Xeon, 70, 72, 74, 78–79, 81–82, 512, 561
- Processor stalls, 89
- Programming model
 ILP-64, 85
 LP-64, 85
- PVFS
 See File system, parallel virtual file system
- PXE
 See Booting, pre-execution environment
- R**
- RAID
 See Redundant array of independent disks
- RDMA
 See Network, protocol, remote direct memory access

- Red Hat package manager, 184, 191, 300, 339
 - binary package, 303–304
 - finding packages, 455
 - installing packages, 343, 355
 - OpenMOSIX packages, 503
 - open-source cluster application resources, 480
 - package database, 302, 499
 - package repository, 575
 - source package, 234, 300, 303
 - SystemImager packages, 354
- Reduced instruction set computing, 9, 80, 158
- Redundant array of independent disks, 71–72, 78, 93, 129, 175, 188
 - mode 0, 188
 - mode 1, 188, 229, 496
 - mode 4, 188
 - mode 5, 129, 188–189, 398–399
 - multipath, 195
 - software, 188–189
 - chunks, 189
 - creating, 192, 197
 - file system stride for mode 5, 211
 - kickstart example, 340
 - mdadm, 200
 - mirroring system partitions, 193
 - mode 1, 191
 - mode 1 example, 198
 - mode 5, 189
 - module dependencies, 191
 - parity calculation, 190
 - querying array members, 199
 - raidtab, 189
 - raidtools, 197
 - recovery, 196, 201
 - spare disk, 196
 - status, 199
 - superblock, 195
 - software devices, 188
- Remote procedure call, 11
- Request for comment, 247
- Resource management system, 437–438, 445, 448
 - Quadrics, 306
 - Resource management system
 - simple Linux utility for, 438
- Resources
 - multitiered compute, 54
- RISC
 - See* Reduced instruction set computing
- RMON
 - See* Network, protocol, remote monitoring
- RPM
 - See* Red Hat package manager
- S**
- SAN
 - See* Storage area network
- SDP
 - See* Network, protocol, Socket Direct Protocol
- Search for extraterrestrial intelligence, 16
 - SETI@home, 52
- Secure shell, 119
 - agent
 - environment variables, 291
 - local port forwarding, 297
 - remote port forwarding, 299
 - RSA and DSA, 290
 - ssh-agent, 291
 - StrictHostKeyChecking, 293
 - X Windows forwarding, 297
- Service
 - fall-back, 224
 - fan-out, 245
 - interference, 222
 - redundant, 223
- Services
 - enabling, 243
 - infrastructure, 37
- SETI
 - See* Search for extraterrestrial intelligence
- SGID
 - See* Authorization, directories, set group identifier

- SI
See Installation, *SystemImager*
- Single point of failure, 57
- Small computer system interconnect, 5
- SMB
See Network, protocol, server message block
- SMP
See System, symmetric multiprocessor
- SNTP
See Network, protocol, Simple Network Time Protocol
- Software
comprehensive Perl archive network, 354
finding
 RPM packages, 354
free, 161
open-source, xxxii, xxxviii, 165–166, 379
 churn, 164
 development, 4, 161, 164, 475
 license, 162–164, 166, 170, 577
 support of, 173
- Software scaling issues
 example of, 39
- Software stack, 38, 134
- Solution, definition of, 13
- SPEC
See Standard Performance Evaluation Corporation
- SRP
See Network, protocol, SCSI remote DMA protocol
- SSH
See Secure shell
- SSL
See Network, secure socket layer
- Standard Performance Evaluation Corporation, 78–79, 81
- Storage area network, 56–57, 59, 61, 92, 130, 134, 401–402, 419, 427–428, 535, 550
- STP
See Network, protocol, Spanning Tree Protocol
- Switch, keyboard–video–mouse, 121, 232
- System
 commercial off the shelf, 52
 symmetric multiprocessor, 6
 uniprocessor, 5
 white box, 522
- System administration
 automation of, 28
- System logging daemon
 log file, 254
- System management
 out of band, 120
- T**
- TCP
See Network, protocol, Transmission Control Protocol
- TCP/IP
See Network, protocol, Transmission Control Protocol/Internet Protocol
- TFTP
See Network, protocol, trivial file transfer protocol
- TGT
See Kerberos, ticket granting ticket
- 3TCS
See Three-tier client–server
- Three-tier client–server, 15
- TLB
See Translation lookaside buffer
- Translation lookaside buffer, 83
- U**
- Universal serial bus, 5, 70, 121, 177, 180, 206, 233, 549
- Universally unique identifier, 189, 200, 202, 208
- USB
See Universal serial bus
- User direct-access programming library, 150



UUID

See Universally unique identifier

V

Version control

concurrent version system, 226

revision control system, 226

VFS

See Virtual file system layer

Virtual address space, 83

W

World Wide Web, xxxviii, 233

WWW

See World Wide Web

X

XML

See Extensible mark-up language

