




Linux Guide

Daniele Raffo

14th edition July 2026

Foreword

This guide stems from the notes I have been taking while studying, teaching, and working with Linux.

It contains information about standards and tools for Linux system administration, as well as a good amount of topics from the certification exams LPIC-1 (Linux Professional Institute Certification level 1), LPIC-2, RHCSA (Red Hat Certified System Administrator), RHCE (Red Hat Certified Engineer), and CEH (Certified Ethical Hacker). Unless otherwise specified, the shell of reference is Bash.

This is an independent publication and is not affiliated with LPI, Red Hat, EC-Council, or any other organization. You are welcome to share this guide, in part or in full, as long as it remains unmodified and is distributed free of charge.

This document was composed with Apache OpenOffice.

Linux is a registered trademark of Linus Torvalds. The Tux penguin logo was created by Larry Ewing.

The Kubernetes images and code examples are © 2026 The Kubernetes Authors (link in the Bibliography below), used under CC BY 4.0 license.

Happy Linux hacking!

Version history

1 st edition	May 2013	9 th edition	January 2021
2 nd edition	September 2014	10 th edition	March 2022
3 rd edition	July 2015	11 th edition	March 2023
4 th edition	June 2016	12 th edition	March 2024
5 th edition	September 2017	13 th edition	January 2025
6 th edition	August 2018	14 th edition	July 2026
7 th edition	May 2019		
8 th edition	January 2020		

Bibliography and resources

- Evi Nemeth et al., *UNIX and Linux System Administration Handbook*, O'Reilly
- Rebecca Thomas et al., *Advanced Programmer's Guide to Unix System V*, McGraw-Hill
- Christoph Braun, *Unix System Security Essentials*, Addison-Wesley
- Jean-Marie Rifflet, *La programmation sous Unix*, Ediscience
- Mendel Cooper, *Advanced Bash-Scripting Guide*, <http://tldp.org/LDP/abs/html>
- Ellen Siever et al., *Linux in a Nutshell*, O'Reilly, <http://archive.oreilly.com/linux/cmd>
- Bruce Barnett, *The Grymoire*, <http://www.grymoire.com/Unix>
- Colin Barschel, *Unix Toolbox*, <http://cb.vu/unixtoolbox.xhtml>
- Adam Haeder et al., *LPI Linux Certification in a Nutshell*, O'Reilly
- Heinrich W. Klöpping et al., *The LPIC-2 Exam Prep*, <http://lpic2.unix.nl>
- Michael Jang, *RHCSA/RHCE Red Hat Linux Certification Study Guide*, McGraw-Hill
- Asghar Ghori, *RHCSA & RHCE RHEL 7: Training and Exam Preparation Guide*, Lightning Source Inc.
- Linus Torvalds' Linux documentation, <https://github.com/torvalds/linux/tree/master/Documentation>
- The Linux Documentation Project guides, <https://www.tldp.org/guides.html>
- RHEL manuals, https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux
- The Linux Kernel Archives, <https://www.kernel.org>
- Linux man pages, <https://www.kernel.org/doc/man-pages>
- Linux man pages, <https://linux.die.net/man>
- A-Z index of Bash command line, <http://ss64.com/bash>
- GNU software manuals, <http://www.gnu.org/manual>
- Shell command line snippets, <http://www.commandlinefu.com>
- Bash command line snippets, <http://www.bashoneliners.com>
- RAM management in Linux, <http://www.linuxatemyram.com>
- Linux performance, <http://www.brendangregg.com/linuxperf.html>
- Bash pitfalls, <http://mywiki.woledge.org/BashPitfalls>
- UNIX & Linux forum, <https://unix.stackexchange.com>
- Explanation for any command and its arguments, <https://explainshell.com>
- Install instructions for any command, <https://command-not-found.com>
- Shell script checker for bugs and errors, <https://www.shellcheck.net>
- Regular expressions tester and cheat sheet, <http://www.regextester.com>
- Cron schedule expression editor, <https://crontab.guru>
- Hands-on lab for Linux troubleshooting exercises, <https://sadservers.com>
- Kubernetes tutorials and documentation, <https://kubernetes.io/docs/tutorials>
- MicroK8s tutorials, <https://microk8s.io/tutorials>
- Linux distributions, <https://distrowatch.com>

Index

LVM.....	1	UID and GID.....	101	DNS - zone file.....	201
LVM - commands.....	2	sudo.....	102	DNS - Resource Records.....	202
System boot.....	3	Terminal.....	103	HTTP response codes.....	203
UEFI.....	4	Terminal - session.....	104	Apache.....	204
SysV startup sequence.....	5	Messaging.....	105	Apache - server configuration.....	205
Login.....	6	cron.....	106	Apache - main configuration.....	206
Runlevels and targets.....	7	at.....	107	Apache - virtual hosts.....	207
Shutdown.....	8	Other schedulers.....	108	Apache - authorization.....	208
SysV service management.....	9	Compilers and debuggers.....	109	Apache - SSL/TLS.....	209
Systemd service management.....	10	Math and database utilities.....	110	Apache - proxy.....	210
/etc/inittab.....	11	Image, audio, and video utilities.....	111	Other web servers.....	211
Filesystem Hierarchy Standard.....	12	Miscellaneous utilities.....	112	Tomcat.....	212
Partitions.....	13	X Window utilities.....	113	Samba - server.....	213
Filesystem.....	14	Linux distributions - part 1.....	114	Samba - client.....	214
fsck.....	15	Linux distributions - part 2.....	115	Samba - global configuration.....	215
mount.....	16	Linux distributions - part 3.....	116	Samba - share configuration.....	216
Partition and filesystem types.....	17	Localization.....	117	Samba - access configuration and macros.....	217
Swap.....	18	System time.....	118	Samba - setup.....	218
/etc/fstab.....	19	NTP.....	119	NFS.....	219
Filesystem operations - part 1.....	20	syslog.....	120	NFS - export table.....	220
Filesystem operations - part 2.....	21	E-mail.....	121	NFS - setup.....	221
inotify and inotifywait.....	22	SMTP.....	122	iSCSI.....	222
XFS, ReiserFS, and CD-ROM filesystems.....	23	Sendmail.....	123	iSCSI - setup.....	223
AutoFS.....	24	Exim.....	124	DHCP.....	224
RAID.....	25	Postfix.....	125	DHCP - message types.....	225
Bootloader.....	26	Postfix - configuration.....	126	PAM.....	226
GRUB 2 - configuration.....	27	Procmail.....	127	LDAP.....	227
GRUB 2 - operations.....	28	Courier - POP configuration.....	128	LDAP - commands.....	228
GRUB Legacy - configuration.....	29	Courier - IMAP configuration.....	129	OpenLDAP.....	229
GRUB Legacy - shell commands.....	30	Dovecot.....	130	389 Directory Server.....	230
dpkg.....	31	Dovecot - mailbox configuration.....	131	SELinux.....	231
apt.....	32	Dovecot - POP and IMAP configuration.....	132	SELinux - semanage.....	232
rpm.....	33	Dovecot - authentication.....	133	SELinux - AVC violations.....	233
yum.....	34	FTP.....	134	SELinux - AVC message format.....	234
dnf.....	35	vsftpd.....	135	auditd.....	235
yum repositories.....	36	CUPS.....	136	Tor.....	236
Other package managers.....	37	IP addressing.....	137	Kickstart.....	237
Backup and data recovery.....	38	Subnetting.....	138	Red Hat Satellite 6.....	238
Tape libraries.....	39	TCP/IP.....	139	KVM.....	239
Archive formats.....	40	Wireless network - standards.....	140	QEMU.....	240
Documentation.....	41	Wireless network - security.....	141	Multipass.....	241
Text filters - part 1.....	42	Network - services.....	142	Docker.....	242
Text filters - part 2.....	43	Network - configuration commands.....	143	Docker - commands.....	243
Text filters - part 3.....	44	Network - configuration files part 1.....	144	Kubernetes - part 1.....	244
File encodings and formats.....	45	Network - configuration files part 2.....	145	Kubernetes - part 2.....	245
Regular expressions.....	46	Network - tools part 1.....	146	Kubernetes - commands.....	246
File management.....	47	Network - tools part 2.....	147	Kubernetes - examples.....	247
Directory management.....	48	Network - wireless tools.....	148	Kubernetes - distributions.....	248
File status.....	49	Network - monitoring.....	149	Cloud computing.....	249
Standard I/O streams.....	50	nmcli.....	150	Kerberos 5.....	250
Standard I/O streams - commands.....	51	nmcli - network teaming and bridging.....	151	Kerberos 5 - commands.....	251
read and echo.....	52	nmcli - network bonding.....	152	RADIUS.....	252
Processes.....	53	nmap.....	153	NSS.....	253
Jobs.....	54	nmap - options part 1.....	154	SSSD.....	254
Signals.....	55	nmap - options part 2.....	155	Identity Management.....	255
kill.....	56	tcpdump.....	156	Dogtag.....	256
Resource monitoring.....	57	netcat.....	157	Git - part 1.....	257
vmstat.....	58	hping3.....	158	Git - part 2.....	258
free.....	59	TCP Wrapper.....	159	Vagrant.....	259
PCP, rasdaemon, and sar.....	60	Routing.....	160	Ceph.....	260
File permissions.....	61	Firewall.....	161	Puppet.....	261
File types.....	62	iptables - rules.....	162	Puppet - example.....	262
File attributes.....	63	iptables - NAT routing.....	163	Puppet - syntax.....	263
ACLs.....	64	firewallld.....	164	Ansible.....	264
Links.....	65	firewallld - rules.....	165	Ansible - playbook example part 1.....	265
Find files and commands.....	66	SSH.....	166	Ansible - playbook example part 2.....	266
Shell usage.....	67	SSH - tools.....	167	Ansible - playbook example part 3.....	267
Shell variables.....	68	SSH - operations.....	168	Ansible Tower.....	268
Shell variables - operations.....	69	SSH - configuration.....	169	HTML 4.01 - components.....	269
Shell mechanics.....	70	X.509.....	170	HTML 4.01 - text.....	270
Shell options.....	71	OpenSSL.....	171	HTML 4.01 - images.....	271
Shell scripting.....	72	OpenSSL - commands.....	172	HTML 4.01 - tables.....	272
getopts.....	73	Cryptography.....	173	7-bit ASCII table.....	273
System information.....	74	Ciphers.....	174	Information Security.....	274
Command execution.....	75	Hash functions.....	175	Metasploit.....	275
Command execution - helper tools.....	76	Authentication systems.....	176	Metasploit - Meterpreter.....	276
Tests.....	77	PGP.....	177	Aircrack-ng - main commands.....	277
Operations.....	78	LUKS.....	178	Aircrack-ng - other commands.....	278
Flow control.....	79	OpenVPN.....	179	Aircrack-ng - recipes.....	279
Text processors.....	80	Auditing.....	180	Firmware Mod Kit.....	280
less.....	81	Fail2Ban.....	181	ClamAV.....	281
Vi - commands.....	82	Key bindings - terminal.....	182	MDE.....	282
Vi - options.....	83	Key bindings - X Window.....	183	Security tools - network.....	283
SQL.....	84	Kernel.....	184	Security tools - wireless.....	284
SQL SELECT.....	85	Kernel - operations.....	185	Security tools - applications.....	285
SQL JOIN.....	86	Kernel - information.....	186	Security tools - web services.....	286
MySQL.....	87	Kernel - compile and patching.....	187	Security tools - passwords.....	287
MySQL - tools.....	88	Kernel - modules.....	188	Security tools - honeypots.....	288
MySQL - commands.....	89	Kernel - logs.....	189	Security tools - misc.....	289
MySQL - status.....	90	/proc.....	190	Attacks - generic.....	290
MySQL - recipes.....	91	/dev - hardware devices.....	191	Attacks - nontechnical.....	291
MySQL - operations.....	92	/dev - virtual devices.....	192	Attacks - DoS.....	292
PostgreSQL.....	93	udev.....	193	Attacks - TCP/IP.....	293
X Window.....	94	System recovery.....	194	Attacks - DNS.....	294
X Window - tools.....	95	System recovery - RHEL password.....	195	Attacks - network.....	295
X Window - keyboard and fonts.....	96	System recovery - RHEL boot.....	196	Attacks - Wi-Fi.....	296
X Window - keysym codes.....	97	System recovery - chmod.....	197	Attacks - Bluetooth.....	297
/etc/passwd.....	98	DNS.....	198	Attacks - hijacking.....	298
User management.....	99	DNSSEC.....	199	Attacks - web services.....	299
Group management.....	100	DNS - configuration.....	200	Attacks - web applications.....	300

Attacks - web.....	301	Attacks - hardware.....	305	Security countermeasures - firewall.....	309
Attacks - XML.....	302	Attacks - cryptography.....	306	Security countermeasures - IDS.....	310
Attacks - SQL.....	303	Malware - virus.....	307	Security countermeasures - WIDS.....	311
Attacks - applications.....	304	Malware - misc.....	308	Security countermeasures - honeypot.....	312

Index of commands

7z.....	40	bbharvest.....	285	crc32.....	49	e2fsck.....	15
a2ps.....	136	bbtrans.....	285	createdb.....	93	e2label.....	14
ac.....	180	bc.....	110	createrepo.....	33	eddefrag.....	21
acct.....	180	bcrypt.....	173	crontab.....	106	easside-ng.....	278
addnupghome.....	177	beep.....	111	crunch.....	287	echo.....	52
addgroup.....	100	bg.....	54	crushtool.....	260	ed.....	80
adduser.....	99	bing.....	146	cryptsetup.....	178	efibootdump.....	4
afetch.....	74	binwalk.....	285	csh.....	67	efibootmgr.....	4
afl-fuzz.....	285	blkdiscard.....	14	csvclean.....	110	efivar.....	4
age.....	173	blkid.....	14	csvcut.....	110	egrep.....	44
aide.....	289	blockdev.....	14	csvformat.....	110	eject.....	16
airbase-ng.....	278	bluediving.....	284	csvgrep.....	110	elif.....	79
aircrack-ng.....	277	bluescan.....	284	csvjoin.....	110	else.....	79
airdecap-ng.....	278	bluesniff.....	284	csvjson.....	110	emacs.....	80
airdecloak-ng.....	278	boxes.....	112	csvlook.....	110	env.....	68
airdriver-ng.....	278	brctl.....	151	csvpy.....	110	envsubst.....	68
airdrop-ng.....	278	break.....	79	csvsort.....	110	esac.....	79
aireplay-ng.....	277	brew.....	37	csvsql.....	110	ethhtool.....	143
airfart.....	284	bts Scanner.....	284	csvstack.....	110	ettercap.....	283
airgraph-ng.....	278	buddy-ng.....	278	csvstat.....	110	evtest.....	186
airmon-ng.....	278	bunzip2.....	40	cups-config.....	136	exicyclog.....	124
airodump-ng.....	277	burp.....	38, 286	cupsaccept.....	136	exiftool.....	111
airolib-ng.....	278	busybox.....	76	cupsaddsmb.....	136	exigrep.....	124
airoscrip-t-ng.....	278	bzcat.....	40	cupsdisable.....	136	exim4.....	124
airserv-ng.....	278	bzip2.....	40	cupsenable.....	136	exinext.....	124
airsnort.....	284	CA.pl.....	171	cupsreject.....	136	exit.....	72
airtun-ng.....	278	cal.....	112	cupstestppd.....	136	exiv2.....	111
alacrity.....	113	case.....	79	curl.....	147	expand.....	42
alias.....	67	cat.....	42	cut.....	43	expect.....	76
alien.....	31	catdoc.....	45	daff.....	110	export.....	68
alpine.....	121	catman.....	41	dash.....	67	exportfs.....	219
amap.....	285	cc.....	109	datamash.....	110	expr.....	77
amapcrap.....	285	ccr.....	173	date.....	118	ext3grep.....	38
ansible.....	264	ccrypt.....	173	dbus-monitor.....	185	ext4magic.....	38
ansible-lint.....	264	cd.....	48	dbx.....	109	extundelete.....	38
ansible-playbook.....	264	cdrwtool.....	23	dc.....	110	factor.....	261
ansible-pull.....	264	ceph.....	260	dd.....	38	factor.....	110
antiword.....	45	ceph-authtool.....	260	ddd.....	109	fail2ban-client.....	181
apache2ctl.....	204	ceph-clsinf-o.....	260	ddosim.....	285	faillog.....	180
apachectl.....	204	ceph-conf.....	260	ddrescue.....	38	fallocate.....	47
aplay.....	111	ceph-mds.....	260	debsecan.....	289	false.....	72
applygnupgdefaults.....	177	ceph-mgr.....	260	debugfs.....	21	fastfetch.....	74
apropos.....	41	ceph-mon.....	260	debugreiserfs.....	23	fat.....	285
apt.....	32	ceph-osd.....	260	delgroup.....	100	fatrat.....	285
apt-cache.....	32	ceph-run.....	260	deluser.....	99	fc.....	67
apt-cdrom.....	32	ceph-volume.....	260	depmod.....	188	fc-cache.....	96
apt-file.....	32	certwatch.....	171	df.....	20	fc-list.....	96
apt-get.....	32	cfdisk.....	13	dhclient.....	143	fc-scan.....	96
apt-key.....	32	chac1.....	64	dhcpcd.....	143	feron.....	108
aptitude.....	32	chage.....	99	dhcpcdump.....	149	fdisk.....	13
arachni_web.....	286	chattr.....	63	dhcpi-g.....	283	fdupes.....	47
arecord.....	111	chcon.....	232	dhcpi-starv.....	283	fern-wifi-cracker.....	284
aria2c.....	147	checkmodule.....	233	dialog.....	72	feroxbuster.....	286
arp.....	143	chfn.....	99	dict.....	112	ferret.....	286
arp-scan.....	149	chgrp.....	61	diff.....	44	ffuf.....	286
arpoison.....	283	chkconfig.....	9	diff-so-fancy.....	44	fg.....	54
arpon.....	283	chkrootkit.....	289	dig.....	146	fi.....	79
arpspoof.....	283	chmod.....	61	dir.....	48	fiddler.....	286
arpstraw.....	283	chown.....	61	dirbuster.....	286	figlet.....	112
arpwatch.....	283	chpasswd.....	99	dirmngr.....	177	file.....	49
artillery.....	288	chromium-browser.....	113	dirmngr-client.....	177	filefrag.....	21
ascii.....	273	chronyc.....	119	dirname.....	48	find.....	66
ash.....	67	chronyd.....	119	dirs.....	48	findfs.....	14
aspell.....	112	chroot.....	20	dirsearch.....	286	findmnt.....	16
at.....	107	chsh.....	99	disown.....	54	finger.....	101
atd.....	107	chvt.....	103	dmesg.....	189	firewall-cmd.....	164
atop.....	57	citellus.py.....	74	dmidecode.....	186	firewall-config.....	164
atq.....	107	ckbcomp.....	96	dmsetup.....	2	firewallld.....	161
atrm.....	107	clamconf.....	281	dnf.....	35	fish.....	67
audit2allow.....	233	clamd.....	281	dnsdomainname.....	143	flatpak.....	37
audit2why.....	233	clamonnacc.....	281	dnssec-keygen.....	199	flock.....	49
auditctl.....	235	clamscan.....	281	dnssec-signzone.....	199	fmt.....	42
auditd.....	235	clang.....	109	dnsspoof.....	283	fold.....	42
aureport.....	235	clear.....	103	dnswalk.....	198	for.....	79
ausearch.....	235	cloc.....	112	do.....	79	foreman-debug.....	238
authconfig.....	229	cmdtest.....	76	doas.....	102	foreman-maintain.....	238
authconfig-gtk.....	229	cmp.....	45	docker.....	243	foreman-rake.....	238
authselect.....	229	collected.....	57	dog.....	289	forkstat.....	56
autossh.....	166	column.....	42	domainname.....	143	formail.....	127
away.....	103	comm.....	44	done.....	79	fortune.....	112
awk.....	43	command.....	66	dpkg.....	31	fping.....	146
awx.....	268	comp-gen.....	66	dpkg-reconfigure.....	31	fragrouter.....	283
awx-cli.....	268	conky.....	57	dracut.....	185	free.....	59
azcmagent.....	282	continue.....	79	dscconf.....	230	freeradius.....	252
balbuzard.....	285	cool-retro-term.....	113	dsctl.....	230	freeramdisk.....	185
banner.....	112	corefreq-cli.....	57	dselect.....	31	freshclam.....	281
bannergrab.....	283	corefreqd.....	57	dsidm.....	230	fsck.....	15
base64.....	45	cowpatty.....	284	dsniff.....	287	fsck.cramfs.....	15
bash.....	67	cowrie.....	288	du.....	20	fsck.ext2.....	15
bastille.....	289	cowsay.....	112	duf.....	20	fsck.ext3.....	15
bat.....	42	cp.....	47	dumpe2fs.....	21	fsck.ext4.....	15
batch.....	108	cpio.....	40	dzen.....	113	fsck.minix.....	15
bbcrack.....	285	cpufetch.....	186	e2freefrag.....	21	fsck.ms-dos.....	15

fsck.vfat.....	15	hping3.....	158	krb5-config.....	251	mailstats.....	123
fsfreeze.....	14	htc.....	283	krb5kdc.....	251	mailx.....	121
fsv.....	113	htop.....	57	ksh.....	67	make.....	109
ftp.....	134	htpasswd.....	204	kstats.....	278	makealiases.....	128
ftpd.....	134	hts.....	283	ksu.....	251	makeivs.....	278
function.....	72	httpbee.....	286	ksvalidator.....	237	makemap.....	123
fuser.....	49	httpd.....	204	ksverdiff.....	237	maltego.....	289
fx.....	45	httpprint.....	286	kswitch.....	251	man.....	41
fzf.....	66	hwclock.....	118	ktutil.....	251	mandb.....	41
fzy.....	66	hydra.....	287	kubeadm.....	246	mapfile.....	52
g++.....	109	hyfetch.....	74	kubectrl.....	246	mariadb-secure-installation.....	87
gasp.....	109	iconv.....	117	kvno.....	251	masscan.....	283
gawk.....	43	id.....	101	labrea.....	288	matchpathcon.....	232
gcc.....	109	if.....	79	last.....	180	mcp.....	47
gdb.....	109	ifconfig.....	143	lastb.....	180	md5pass.....	173
gdm.....	95	ifdown.....	143	lastcomm.....	180	md5sum.....	173
gedit.....	80	ifne.....	51	latex.....	112	mdadm.....	25
genkey.....	171	ifup.....	143	ld.....	109	mdatp.....	282
genpmk.....	284	ike-scan.....	285	ldapadd.....	228	medusa.....	287
getenforce.....	231	in2csv.....	110	ldapdelete.....	228	mesg.....	105
getent.....	253	incrontab.....	22	ldapmodify.....	228	microk8s.....	248
getfaci.....	64	inetsim.....	147	ldappasswd.....	228	minikube.....	248
getopts.....	73	info.....	41	ldapsearch.....	228	mitmf.....	283
getsebool.....	232	init.....	7	ldconfig.....	185	mkdir.....	48
gfetxh.....	258	innotop.....	88	ldd.....	185	mkdosfs.....	14
gh.....	258	inotifywait.....	22	less.....	81	mke2fs.....	14
ghostwriter.....	113	inotifywatch.....	22	let.....	68	mke3fs.....	14
git.....	257	insmod.....	188	lftp.....	134	mkfifo.....	50
gkrellm.....	57	inspy.....	289	lighttpd.....	211	mkfontdir.....	96
gksu.....	102	inxi.....	57	l1c.....	109	mkfs.....	14
gksudo.....	102	iodine.....	283	lldb.....	109	mkfs.ext2.....	14
glab.....	258	iodined.....	283	ln.....	65	mkfs.ext3.....	14
gnome-cups-manager.....	136	iostat.....	58	loadkeys.....	96	mkfs.ext4.....	14
gnome-device-manager.....	193	iotop.....	57	locale.....	117	mkfs.jfs.....	14
gnome-disks.....	13	ip.....	143	locale-gen.....	117	mkfs.msdoos.....	14
gnome-terminal.....	113	ip6tables.....	161	locate.....	66	mkfs.ntfs.....	14
gnome-tweaks.....	113	ipa.....	255	lockfile.....	127	mkfs.reiserfs.....	14
gnuplot.....	110	ipcalc.....	112	logger.....	120	mkfs.xfs.....	14
goaccess.....	112	ipchains.....	161	logresolve.....	204	mkimapdcert.....	128
gobuster.....	286	ipcs.....	56	logrotate.....	120	mkinitramfs.....	185
google-authenticator.....	112	iperf.....	149	loic.....	283	mkinitrd.....	185
google-chrome.....	113	ipfwadm.....	161	loki.....	283	mkisofs.....	23
gotty.....	112	ipgrab.....	149	lolcat.....	112	mknod.....	20
gpated.....	13	iptables.....	161	look.....	44	mkntfs.....	14
gpaswd.....	100	iptables-restore.....	161	losetup.....	20	mkpop3dcert.....	128
gpg.....	177	iptables-save.....	161	lp.....	136	mkreiserfs.....	14
gpg-check-pattern.....	177	iptraf.....	149	lpadmin.....	136	mkswap.....	18
gpg-connect-agent.....	177	iptraf-ng.....	149	lpc.....	136	mktemp.....	47
gpg-preset-passphrase.....	177	iscsiadm.....	222	lpq.....	136	mkudffs.....	23
gpgconf.....	177	ISR-sqlget.....	285	lprm.....	136	mln.....	47
gpgtar.....	177	ivstools.....	278	lpstat.....	136	mlr.....	110
gpgv.....	177	iw.....	148	ls.....	48	mmv.....	47
grep.....	44	iwconfig.....	148	lsattr.....	63	modinfo.....	188
grepcidr.....	112	iwgetid.....	148	lsb_release.....	74	modprobe.....	188
groupadd.....	100	iwlist.....	148	lsblk.....	20	mole.....	285
groupdel.....	100	java.....	212	lscpu.....	186	monmaptool.....	260
groupmod.....	100	jobber.....	108	lsdev.....	186	more.....	80
grub.....	28	jobs.....	54	lsdiff.....	109	most.....	80
grub-emu.....	28	john.....	287	lsdvd.....	111	mount.....	16
grub-install.....	28	join.....	42	lshal.....	193	mount.cifs.....	214
grub2-editenv.....	28	jot.....	110	lshw.....	186	mountpoint.....	16
grub2-mkconfig.....	28	journalctl.....	189	lsinitramfs.....	185	mpage.....	136
grub2-set-default.....	28	jp.....	110	ls1k.....	49	mpd.....	111
grubby.....	28	jp2a.....	112	lslocks.....	49	mpstat.....	58
gunzip.....	40	jq.....	45	lsmbos.....	121	msfconsole.....	275
gv.....	136	json_reformat.....	45	lsmod.....	188	msfencode.....	275
gvim.....	80	json_verify.....	45	lsosf.....	49	msfpayload.....	275
gzip.....	40	jsql.....	285	lspci.....	186	msfupdate.....	275
hackrf_info.....	284	JustMetadata.....	289	lspcmcia.....	186	msfvenom.....	275
hadolint.....	243	k5srvutil.....	251	lsscsi.....	20	mt.....	39
hald.....	193	k9s.....	246	lsusb.....	186	mtr.....	146
halt.....	8	kadmin.....	251	lt.....	147	mtx.....	39
hammer.....	238	kadmin.local.....	251	lvchange.....	2	multipass.....	241
hamster.....	286	kadmind.....	251	lvcreate.....	2	multipath.....	20
hashcat.....	287	katoolin.....	32	lvdisplay.....	2	multitail.....	43
haveged.....	112	kbdcontrol.....	96	lvextend.....	2	mutt.....	121
hcidump.....	148	kdb5_ldap_util.....	251	lvm.....	2	mv.....	47
hd.....	45	kdb5_util.....	251	lvmconfig.....	2	mysql.....	87
hdparm.....	21	kdestroy.....	251	lvmdiskscan.....	2	mysql_install_db.....	87
head.....	43	kdm.....	95	lvmsar.....	2	mysql_secure_installation.....	87
help.....	41	kernelld.....	188	lvreduce.....	2	mysql_upgrade.....	87
hexdump.....	45	kexec.....	185	lvremove.....	2	mysqlcheck.....	87
hexedit.....	45	keyrand.....	171	lvresize.....	2	mysqld.....	87
hexeditor.....	45	kill.....	56	lvs.....	2	mysqld_safe.....	87
hfetch.....	74	killall.....	56	lvscan.....	2	mysqldump.....	87
history.....	67	kinit.....	251	lynis.....	289	mysqlreport.....	88
hoic.....	283	kismet.....	284	lzcac.....	40	mysqlslap.....	88
honeyd.....	288	kitty.....	113	lzma.....	40	mysq1tuner.pl.....	88
honeyp.....	288	klist.....	251	m4.....	123	mytop.....	88
honeytrap.....	288	klogd.....	120	macchanger.....	283	named.....	198
horsepill.....	285	kmod.....	188	macof.....	283	name1.....	49
host.....	146	kojoney.....	288	mad.....	47	nano.....	80
hostapd.....	148	kpasswd.....	251	magick.....	111	nasm.....	109
hostname.....	143	kprop.....	251	mail.....	121	nbtscan.....	283
hostnamectl.....	143	kpropd.....	251	mailg.....	123	nc.....	157
hoststat.....	123	kproplog.....	251	mailstat.....	127	ncat.....	157

ncdu.....	20	pirut.....	33	gemu-system-s390x.....	240	setfacl.....	64
ncmpc.....	111	pkexec.....	102	gemu-system-sh4.....	240	setsebool.....	232
ncmpcpp.....	111	pkill.....	56	gemu-system-sparc.....	240	sf.....	286
ncrack.....	287	pktgen.....	147	gemu-system-tricore.....	240	sfdisk.....	13
ndc.....	198	pmap.....	53	gemu-system-unicore32.....	240	sftp.....	167
ndn.....	112	pmc.....	119	gemu-system-x86_64.....	240	sh.....	67
needs-restarting.....	34	pmddmptext.....	60	gemu-system-xtensa.....	240	shalsum.....	173
neofetch.....	74	pmevent.....	60	gemu-x86_64.....	240	sha224sum.....	173
neowofetch.....	74	pminfo.....	60	radamsa.....	285	sha256sum.....	173
nerdfetch.....	74	pmrep.....	60	radclient.....	252	sha384sum.....	173
nessusd.....	289	pmstat.....	60	radiusd.....	252	sha512sum.....	173
net.....	213	pmval.....	60	rados.....	260	shasum.....	173
netcat.....	157	popd.....	48	radtest.....	252	shc.....	109
nethogs.....	149	postalias.....	125	rar.....	40	shellcheck.....	72
netperf.....	149	postconf.....	125	ras-mc-ctl.....	60	shopt.....	71
netserver.....	149	postfix.....	125	rasdaemon.....	60	showkey.....	273
netstat.....	149	postgresql-setup.....	93	rcclone.....	38	showmount.....	219
network.....	143	postmap.....	125	rcrack.....	287	shred.....	47
networking.....	143	postqueue.....	125	rdisc.....	160	shuf.....	42
newaliases.....	123	postsuper.....	125	read.....	52	shutdown.....	8
nexphisher.....	289	poweroff.....	8	readarray.....	52	sk.....	66
nextfetch.....	74	powertop.....	57	readelf.....	109	skill.....	56
nfsstat.....	219	pr.....	42	readonly.....	68	skipfish.....	286
nft.....	161	praliases.....	123	realpath.....	48	sl.....	112
nginx.....	211	procmail.....	127	reboot.....	8	slapadd.....	229
ngrep.....	149	ps.....	53	recode.....	117	slapcat.....	229
nice.....	53	ps2pdf.....	136	recon-ng.....	289	slapindex.....	229
nikto.....	286	psad.....	289	redir.....	147	slaptest.....	229
nisdomainname.....	143	psensor.....	57	reiserfstune.....	23	sleep.....	76
nl.....	42	psensor-server.....	57	rename.....	47	slocate.....	66
nload.....	149	pspy.....	285	renice.....	53	slowhttptest.....	286
nmap.....	153	psql.....	93	repoquery.....	34	slowloris.....	286
nmdb.....	213	pssh.....	166	reptyr.....	104	smartctl.....	21
nmblookup.....	213	pstree.....	53	reset.....	103	smartd.....	21
nmcli.....	150	ptp4l.....	119	resize2fs.....	14	smbclient.....	214
nmtui.....	150	pulp-admin-client.....	238	responder.....	283	smbd.....	213
nnn.....	112	pump.....	143	restorecon.....	232	smbmount.....	214
nohup.....	54	puppet.....	261	rev.....	42	smbpasswd.....	213
nologin.....	6	pure-ftpd.....	134	rfcacat.....	284	smbstatus.....	214
notify-send.....	113	pure-ftppwho.....	134	rfcacat-rolljam.py.....	284	snap.....	37
npm.....	37	pure-mrtginfo.....	134	RFCrack.py.....	284	snice.....	53
nscan.....	283	pure-pw.....	134	rfkill.....	148	snoop.....	149
nslookup.....	146	pure-pwconvert.....	134	rkhunter.....	289	snort.....	289
ntfsundelete.....	38	pure-quotacheck.....	134	rlwrap.....	75	socat.....	147
ntop.....	149	pure-statsdecode.....	134	rm.....	47	sort.....	42
ntopng.....	149	pure-uploadsript.....	134	rmdir.....	48	sos.....	74
ntpd.....	119	purgestat.....	123	rmmod.....	188	sosreport.....	74
ntpddate.....	119	pushd.....	48	rnano.....	80	source.....	72
ntpdcc.....	119	putty.....	113	rndc.....	198	sox.....	111
ntpg.....	119	pv.....	51	rndc-confgen.....	199	speaker-test.....	111
ntptrace.....	119	pvchange.....	2	route.....	160	split.....	42
nuclei.....	286	pvck.....	2	rpc.lockd.....	219	sponge.....	51
nvtop.....	57	pvcreeate.....	2	rpc.mountd.....	219	sq.....	177
objdump.....	109	pvddisplay.....	2	rpc.nfsd.....	219	sq-git.....	177
oc.....	249	pvmmove.....	2	rpc.statd.....	219	sql2csv.....	110
od.....	45	pvremove.....	2	rpcinfo.....	219	SQLix.....	286
on_ac_power.....	112	pvresize.....	2	rpl.....	43	sqlmap.....	285
onefetch.....	258	pvs.....	2	rpm.....	33	sqlninja.....	286
openssl.....	171	pvsan.....	2	rpm2cpio.....	33	sqlsus.....	285
openvpn.....	179	pwd.....	48	rsync.....	38	sqv.....	177
ophcrack-cli.....	287	pwgen.....	112	rsyslogd.....	120	ss.....	149
opkg.....	37	pwggen.....	112	rtgen.....	287	sserver.....	251
oscap.....	289	q.....	110	run-parts.....	75	ssh.....	166
osdmapprool.....	260	qalc.....	113	runlevel.....	7	ssh-add.....	166
ossec-hids.....	289	qalculate-gtk.....	113	runuser.....	102	ssh-agent.....	166
p0f.....	283	gemu-alpha.....	240	rusers.....	101	ssh-copy-id.....	166
packetforge-ng.....	278	gemu-arm.....	240	rwho.....	101	ssh-keygen.....	166
packETH.....	147	gemu-armeb.....	240	sa.....	180	ssh-keyscan.....	166
packETHcli.....	147	gemu-i386.....	240	saint.....	289	sshfs.....	167
pacman.....	37	gemu-m68k.....	240	sar.....	60	sshhpot.....	288
pam_tally2.....	180	gemu-microblaze.....	240	scalpel.....	38	sshpass.....	166
pandoc.....	45	gemu-mips.....	240	scapy.....	283	sslstrip.....	286
parallel.....	75	gemu-mips64.....	240	scapy-radio.....	284	sss_cache.....	254
paros.....	286	gemu-mips64el.....	240	sclient.....	251	sssd.....	254
parted.....	13	gemu-mipsel.....	240	scp.....	167	stap.....	186
partprobe.....	13	gemu-mipsn32.....	240	scponly.....	167	stap-prep.....	186
pasfetch.....	74	gemu-mipsn32el.....	240	scrapy.....	286	stap-report.....	186
passwd.....	99	gemu-ppc.....	240	screen.....	104	star.....	40
paste.....	42	gemu-ppc64.....	240	screenfetch.....	74	startx.....	95
patator.....	287	gemu-sh4.....	240	script.....	76	stat.....	49
patch.....	109	gemu-sh4eb.....	240	scrot.....	111	stdbuf.....	51
pcp.....	60	gemu-sparc.....	240	sdparm.....	21	stegdetect.....	289
pentbox.....	289	gemu-sparc32plus.....	240	sealert.....	233	steghide.....	173
pfetch.....	74	gemu-sparc64.....	240	sed.....	43	stegsnow.....	173
pgawk.....	43	gemu-system-alpha.....	240	seinfo.....	233	strace.....	56
pgrep.....	53	gemu-system-arm.....	240	semanage.....	232	stress-ng.....	57
phc2sys.....	119	gemu-system-hppa.....	240	semodule.....	233	strings.....	44
photorec.....	38	gemu-system-i386.....	240	semodule_package.....	233	stty.....	103
pico.....	80	gemu-system-lm32.....	240	sendmail.....	123	stunnel.....	147
pidof.....	53	gemu-system-m68k.....	240	sensors.....	57	su.....	102
pigz.....	40	gemu-system-microblazeel.....	240	sepolicy.....	233	sublist3r.....	286
pine.....	121	gemu-system-mips.....	240	seq.....	110	subscription-manager.....	238
pinentry.....	177	gemu-system-nios2.....	240	service.....	9	sudo.....	102
ping.....	146	gemu-system-or1k.....	240	sestatus.....	231	sudoedit.....	102
pinky.....	101	gemu-system-ppc.....	240	set.....	68	svcrack.....	285
pip.....	37	gemu-system-riscv64.....	240	setenforce.....	231	svcrash.....	285

svmap.....	285	traceroute.....	146	vgcreate.....	2	xautolock.....	95
svreport.....	285	trafgen.....	147	vgdisplay.....	2	xclip.....	95
swar.....	285	traitor.....	285	vgexport.....	2	xdm.....	95
swaks.....	121	trap.....	56	vgextend.....	2	xdmtool.....	95
swapoff.....	18	tree.....	48	vgimport.....	2	xdpyinfo.....	95
swapon.....	18	tripwire.....	289	vgmerge.....	2	xfs.....	96
switchdesk.....	95	trivy.....	243	vgreduce.....	2	xfs_check.....	23
synaptic.....	32	true.....	72	vgremove.....	2	xfs_db.....	23
sync.....	20	truncate.....	47	vgs.....	2	xfs_fsr.....	23
sysbench.....	57	ts.....	75	vgscan.....	2	xfs_growfs.....	23
sysctl.....	190	tset.....	103	vgsplit.....	2	xfs_info.....	23
syslogd.....	120	tty.....	103	vi.....	82	xfs_repair.....	23
sysmon.....	57	tune2fs.....	21	vigr.....	99	xfsdump.....	23
system-config-authentication.....	229	turbostat.....	57	vim.....	80	xfsrestore.....	23
system-config-display.....	95	type.....	66	vimdiff.....	80	xhost.....	95
system-config-firewall.....	161	typeset.....	72	vipw.....	99	xinetd.....	142
system-config-kickstart.....	237	typst.....	112	virsh.....	239	xkbcomp.....	96
system-config-users.....	99	tzconfig.....	118	virt-clone.....	239	xkill.....	56
systemctl.....	10	tzselect.....	118	virt-install.....	239	xmodmap.....	96
systemd-machine-id-setup.....	6	udevadm.....	193	virt-manager.....	239	Xorg.....	95
tac.....	42	udev.....	193	virt-what.....	239	xorgcfg.....	95
tail.....	43	udevmonitor.....	193	virt-who.....	238	xorgconfig.....	95
talk.....	105	udffsck.....	23	visudo.....	102	xosview.....	95
tar.....	40	ufetch.....	74	vlock.....	103	xprop.....	95
targetcli.....	222	ufw.....	161	vmstat.....	58	xrandr.....	95
tcpd.....	147	umask.....	61	vsftpd.....	135	xsel.....	95
tcpdump.....	156	umount.....	16	w.....	101	xset.....	95
tcptrace.....	149	unalias.....	67	w3af.....	286	xsser.....	286
tcptraceroute.....	146	uname.....	186	w3af-console.....	286	xwininfo.....	95
tcsh.....	67	unexpand.....	42	wafninja.....	286	xxd.....	45
teamctl.....	151	unig.....	42	wafw00f.....	286	xz.....	40
teamnl.....	151	uniscan.....	286	wall.....	105	xcats.....	40
tee.....	51	units.....	110	wapiti.....	286	xdm.....	95
telinit.....	7	unlink.....	47	watch.....	75	yasm.....	109
telnet.....	146	unlзма.....	40	watchgnupg.....	177	yersinia.....	283
tesseract.....	112	unrar.....	40	wc.....	42	yes.....	76
test.....	77	unset.....	68	wcalc.....	113	youtube-dl.....	111
testdisk.....	38	until.....	79	webScarab.....	286	ypdomainname.....	143
testparm.....	213	unxz.....	40	websploit.....	289	yum.....	34
textql.....	110	unzstd.....	40	wepattack.....	284	yum-complete-transaction.....	34
tftp.....	134	update-grub.....	28	WEPCrack.....	284	yumdownloader.....	34
then.....	79	update-rc.d.....	9	wesside-ng.....	278	yup.....	34
thttpd.....	211	updatedb.....	66	wget.....	147	zap.....	286
tidy.....	45	uptime.....	57	whatis.....	41	zbarimg.....	111
time.....	75	useradd.....	99	whatweb.....	286	zcat.....	40
timedatectl.....	118	userdel.....	99	whence.....	66	zdump.....	118
timeout.....	75	usermod.....	99	where.....	66	zenity.....	72
tkiptun-ng.....	278	users.....	101	whereis.....	66	zgrep.....	40
tlcr.....	147	usleep.....	76	which.....	66	zless.....	40
tmpwatch.....	49	uucp.....	147	while.....	79	zmap.....	283
tmux.....	104	uudecode.....	45	who.....	101	zmore.....	40
toilet.....	112	uuencode.....	45	whoami.....	101	zsh.....	67
top.....	57	uuidgen.....	112	whoas.....	33	zstd.....	40
tor.....	236	vagrant.....	259	whois.....	146	zstdcat.....	40
tor-router.....	236	valgrind.....	109	wipefs.....	14	zstdmt.....	40
torsocks.....	236	vd.....	110	wpscan.....	286	zz.....	283
touch.....	47	vdir.....	48	write.....	105		72
tower-cli.....	268	vega.....	286	wrudf.....	23		72
tput.....	103	versuck-ng.....	278	wsfuzzer.....	286	[.....	77
tr.....	43	vgchange.....	2	X.....	95		
tracepath.....	146	vgck.....	2	xargs.....	75		

Logical Volume Management (LVM) introduces an abstraction between physical and logical storage, allowing a more versatile use of filesystems. LVM uses the Linux device mapper feature (`/dev/mapper`).

Disks, partitions, and RAID devices are made of **Physical Volumes**, which are grouped into a **Volume Group**. A Volume Group is divided into small fixed-size chunks called Physical Extents, which are mapped 1-to-1 to Logical Extents. Logical Extents are grouped into **Logical Volumes**, on which filesystems are created.

How to create a Logical Volume

1. Add a new disk to the machine
2. `lsblk` Verify that the new disk is recognized e.g. as `/dev/sda`
3. `fdisk /dev/sda` Create a new partition (of type 0x8E = Linux LVM) on the new disk. This is not necessary but recommended, because other OSes might not recognize the LVM header and see the whole unpartitioned disk as empty
4. `pvccreate /dev/sda1` Initialize the Physical Volume to be used with LVM
5. `vgcreate -s 8M myvg /dev/sda1` Create a Volume Group and define the size of Physical Extents to 8 Mb (default value is 4 Mb)
- or 5. `vgextend myvg /dev/sda1` or add the Physical Volume to an existing Volume Group
6. `lvcreate -L 1024M -n mylv myvg` Create a Logical Volume
7. `mkfs -t ext3 /dev/myvg/mylv` Create a filesystem on the Logical Volume
8. `mount /dev/myvg/mylv /mnt/mystuff` Mount the Logical Volume

How to increase the size of a Logical Volume (operation possible only if the underlying filesystem allows it)

1. Add a new disk to the machine, to provide the extra disk space
2. `pvccreate /dev/sdc` Initialize the Physical Volume e.g. `/dev/sdc`
3. `vgextend myvg /dev/sdc` Add the Physical Volume to an existing Volume Group
- or
1. Increase the size of an existing virtual disk (already initialized as a PV, already in a VG, and with LVs allocated on it)
2. `partprobe` Notify the kernel of the new disk size
3. `pvresize /dev/sdc` Accommodate the Physical Volume to the new size
- Then:
4. `lvextend -L 2048M /dev/myvg/mylv` Extend the Logical Volume by 2 Gb
- or 4. `lvresize -L+2048M /dev/myvg/mylv`
- or 4. `lvresize -l+100%FREE /dev/myvg/mylv` or extend the Logical Volume taking all free space
5. `resize2fs /dev/myvg/mylv (ext)` Extend the filesystem.
- `xfs_growfs /dev/myvg/mylv (XFS)` Alternatively, use `lvresize -r` on the previous step

How to reduce the size of a Logical Volume (operation possible only if the underlying filesystem allows it)

1. `resize2fs /dev/myvg/mylv 900M` Shrink the filesystem to 900 Mb
2. `lvreduce -L 900M /dev/myvg/mylv` Shrink the Logical Volume to 900 Mb
- or 2. `lvresize -L 900M /dev/myvg/mylv`

How to snapshot and backup a Logical Volume

1. `lvcreate -s -L 1024M -n mysnap /dev/myvg/mylv` Create the snapshot like a Logical Volume
2. `tar cvzf mysnap.tar.gz mysnap` Backup the snapshot with any backup tool
3. `lvremove /dev/mvvg0/mysnap` Delete the snapshot

PV commands		VG commands		LV commands	
<code>pvs</code>	Report information about Physical Volumes	<code>vgs</code>	Report information about Volume Groups	<code>lvs</code>	Report information about Logical Volumes
<code>pvscan</code>	Scan all disks for Physical Volumes	<code>vgscan</code>	Scan all disks for Volume Groups	<code>lvscan</code>	Scan all disks for Logical Volumes
<code>pvdisplay</code>	Display Physical Volume attributes	<code>vgdisplay</code>	Display Volume Group attributes	<code>lvdisplay</code>	Display Logical Volume attributes
<code>pvck</code>	Check Physical Volume metadata	<code>vgck</code>	Check Volume Group metadata		
<code>pvcreate</code>	Initialize a disk or partition for use with LVM	<code>vgcreate</code>	Create a Volume Group using Physical Volumes	<code>lvcreate</code>	Create a Logical Volume in a Volume Group
<code>pvchange</code>	Change Physical Volume attributes	<code>vgchange</code>	Change Volume Group attributes	<code>lvchange</code>	Change Logical Volume attributes
<code>pvremove</code>	Remove a Physical Volume	<code>vgremove</code>	Remove a Volume Group	<code>lvremove</code>	Remove a Logical Volume
		<code>vgextend</code>	Add a Physical Volume to a Volume Group	<code>lvextend</code>	Increase the size of a Logical Volume
		<code>vgreduce</code>	Remove a Physical Volume from a Volume Group	<code>lvreduce</code>	Reduce the size of a Logical Volume
<code>pvresize</code>	Modify the size of a Physical Volume			<code>lvresize</code>	Modify the size of a Logical Volume
		<code>vgmerge</code>	Merge two Volume Groups		
		<code>vgsplit</code>	Split two Volume Groups		
		<code>vgimport</code>	Import a Volume Group into a system		
		<code>vgexport</code>	Export a Volume Group from a system		
<code>pvmove</code>	Move the Logical Extents on a Physical Volume to wherever there are available Physical Extents (within the Volume Group) and then put the Physical Volume offline				
LVM global commands					
<code>dmsetup command</code>	Perform low-level LVM operations				
<code>lvm command</code>	Perform LVM operations. May also be used as an interactive tool				
<code>lvmsar</code>	LVM system activity reporter. Unsupported on LVM2				
<code>lvmdiskscan</code>	Scan the system for disks and partitions usable by LVM				
<code>lvmconfig</code>	Show the current LVM disk configuration				

`/dev/mapper/vgname-lvname`
`/dev/vgname/lvname`

Mapping of Logical Volumes in the filesystem

`/etc/lvm/archive/`

Directory containing Volume Groups metadata backups

Boot sequence	
POST (Power-On Self Test)	Low-level check of PC hardware.
BIOS (Basic I/O System)	Detection of disks and hardware.
Chain loader GRUB	GRUB stage 1 is loaded from the MBR and executes GRUB stage 2 from filesystem. GRUB prompts the user to select a Linux kernel. The chain loader hands over to the boot sector of the partition on which resides the OS, then executes the corresponding kernel image.
Linux kernel	If the kernel image follows the <code>initrd</code> scheme, the chain loader mounts <code>initrd</code>, an initial ramdisk (typically a compressed ext2 filesystem, with its driver compiled statically into the kernel) to be used as the initial root device during kernel boot; this makes possible to load kernel modules that recognize hard drives hardware and that are hence needed to mount the real root filesystem. Afterwards, the system runs <code>/linuxrc</code> with PID 1. If the kernel image follows the <code>initramfs</code> scheme (available from Linux 2.6.13 onwards) instead, the system loads into memory <code>initramfs</code>, a cpio-compressed image, and unpacks it into an instance of <code>tmpfs</code> in RAM. The kernel is decompressed into memory and executed. The system detects any attached device. The real root filesystem is mounted on <code>/</code> in place of the initial ramdisk.
init	In SysV machines, the kernel executes <code>init</code>, the first process (PID 1). The system tries to execute in order <code>/sbin/init</code>, <code>/etc/init</code>, <code>/bin/init</code>, and <code>/bin/sh</code>; if none of these succeeds, the kernel panics. Then the system loads startup scripts and runlevel scripts. In systemd machines, the kernel executes <code>systemd</code> (PID 1). The system mounts filesystems and swap partitions, and starts low-level services (<code>sysinit.target</code>). The system sets various timers, paths, and sockets (<code>basic.target</code>). The system loads <code>default.target</code>, which is either <code>multi-user.target</code> i.e. console login in text mode or <code>graphical.target</code> i.e. graphical login.
Login	If in text mode, the system calls the <code>getty</code> process, which runs the <code>login</code> command that asks the user for login and password. If in graphical mode, the X Display Manager starts the X Server.

Information about the boot process can be found in the manpages `man 7 boot` and `man 7 bootup`.

Modern systems use **UEFI (Unified Extensible Firmware Interface)** instead of BIOS. UEFI does not use the MBR boot code; it has knowledge of partition table and filesystems, and stores its application files required for launch in an EFI System Partition, usually formatted as FAT32.

After the POST, the system loads the UEFI firmware which initializes the hardware required for booting, then reads its Boot Manager data to determine which UEFI application to launch. The launched UEFI application may then launch another application, e.g. the kernel and initramfs in case of a boot loader like GRUB.

<code>efivar</code>	Manipulate UEFI variables
<code>efibootmgr</code>	Manipulate the UEFI Boot Manager
<code>efibootdump</code>	Display UEFI boot options

Startup sequence	Debian	Red Hat
At startup <code>/sbin/init</code> executes all instructions on <code>/etc/inittab</code> . This script at first switches to the default runlevel...	<code>id:2:initdefault:</code>	<code>id:5:initdefault:</code>
... then it runs the following script (same for all runlevels) which configures peripheral hardware, applies kernel parameters, sets hostname, and provides disks initialization...	<code>/etc/init.d/rcS</code>	<code>/etc/rc.d/rc.sysinit</code> or <code>/etc/rc.sysinit</code>
... and then, for runlevel <i>N</i> , it calls the script <code>/etc/init.d/rc N</code> (i.e. with the runlevel number as parameter) which launches all services and daemons specified in the following startup directories:	<code>/etc/rcN.d/</code>	<code>/etc/rc.d/rcN.d/</code>
The startup directories contain symlinks to the init scripts in <code>/etc/init.d/</code> which are executed in numerical order. Links starting with K are called with argument <code>stop</code>, links starting with S are called with argument <code>start</code>. <pre>lrwxrwxrwx. 1 root root 14 Feb 11 22:32 K88sssd -> ../init.d/sssd lrwxrwxrwx. 1 root root 15 Nov 28 14:50 K89rdisc -> ../init.d/rdisc lrwxrwxrwx. 1 root root 17 Nov 28 15:01 S01sysstat -> ../init.d/sysstat lrwxrwxrwx. 1 root root 18 Nov 28 14:54 S05cgconfig -> ../init.d/cgconfig lrwxrwxrwx. 1 root root 16 Nov 28 14:52 S07iscsid -> ../init.d/iscsid lrwxrwxrwx. 1 root root 18 Nov 28 14:42 S08iptables -> ../init.d/iptables</pre> The last script to be run is <code>S99local -> ../init.d/rc.local</code>; therefore, an easy way to run a specific program upon boot is to call it from this script file.		
<code>/etc/init.d/boot.local</code>	runs only at boot time, not when switching runlevel.	
<code>/etc/init.d/before.local</code> (SUSE)	runs only at boot time, before the scripts in the startup directories.	
<code>/etc/init.d/after.local</code> (SUSE)	runs only at boot time, after the scripts in the startup directories.	
To add or remove services at boot sequence:	<code>update-rc.d service defaults</code> <code>update-rc.d -f service remove</code>	<code>chkconfig --add service</code> <code>chkconfig --del service</code>
When adding or removing a service at boot, startup directories will be updated by creating or deleting symlinks for the default runlevels: <code>K</code> symlinks for runlevels 0 1 6, and <code>S</code> symlinks for runlevels 2 3 4 5. Service will be run via the <code>xinetd</code> super server.		

Linux Standard Base (LSB)
The Linux Standard Base defines a format to specify default values on an init script <code>/etc/init.d/foo</code>: <pre>### BEGIN INIT INFO # Provides: foo # Required-Start: bar # Default-Start: 2 3 4 5 # Default-Stop: 0 1 6 # Description: Service Foo init script ### END INIT INFO</pre> Default runlevels and <code>S/K</code> symlinks values can also be specified as such: <pre># chkconfig: 2345 85 15 # description: Foo service</pre>

<code>/etc/init/start-ttys.conf</code> (Red Hat)	Start the specified number of terminals at startup via <code>getty</code> , which manages physical or virtual terminals (TTYs)												
<code>/etc/sysconfig/init</code> (Red Hat)	Control appearance and functioning of the system during bootup												
<code>/etc/machine-id</code> (Red Hat)	Randomly-generated machine ID. Can be safely regenerated via the command <code>rm -f /etc/machine-id && systemd-machine-id-setup</code>												
<code>/etc/securetty</code>	List of TTYs from which the root user is allowed to login												
<code>/etc/issue</code>	Message printed before the login prompt. Can contain these escape codes: <table> <tr> <td><code>\b</code> Baudrate of line</td><td><code>\o</code> Domain name</td></tr> <tr> <td><code>\d</code> Date</td><td><code>\r</code> OS release number</td></tr> <tr> <td><code>\s</code> System name and OS</td><td><code>\t</code> Time</td></tr> <tr> <td><code>\l</code> Terminal device line</td><td><code>\u</code> Number of users logged in</td></tr> <tr> <td><code>\m</code> Machine architecture identifier</td><td><code>\U</code> "<i>n</i> users" logged in</td></tr> <tr> <td><code>\n</code> Nodename aka hostname</td><td><code>\v</code> OS version and build date</td></tr> </table>	<code>\b</code> Baudrate of line	<code>\o</code> Domain name	<code>\d</code> Date	<code>\r</code> OS release number	<code>\s</code> System name and OS	<code>\t</code> Time	<code>\l</code> Terminal device line	<code>\u</code> Number of users logged in	<code>\m</code> Machine architecture identifier	<code>\U</code> " <i>n</i> users" logged in	<code>\n</code> Nodename aka hostname	<code>\v</code> OS version and build date
<code>\b</code> Baudrate of line	<code>\o</code> Domain name												
<code>\d</code> Date	<code>\r</code> OS release number												
<code>\s</code> System name and OS	<code>\t</code> Time												
<code>\l</code> Terminal device line	<code>\u</code> Number of users logged in												
<code>\m</code> Machine architecture identifier	<code>\U</code> " <i>n</i> users" logged in												
<code>\n</code> Nodename aka hostname	<code>\v</code> OS version and build date												
<code>/etc/issue.net</code>	Message printed before the login prompt on a remote session												
<code>/etc/motd</code>	Message Of The Day, printed after a successful login, but before execution of the login shell												
<code>/etc/nologin</code>	If this file exists, <code>login</code> and <code>sshd</code> deny login to all unprivileged users. Useful when doing system maintenance												
<code>/etc/login.defs</code>	Definition of default values (UID and GID ranges, mail directory, account validity, password encryption method, etc.) for user account creation												
<code>/var/log/secure</code> (Red Hat)	Logfile containing user logins (both successful and failed) and authentication mechanisms												
<code>/var/log/auth.log</code> (Debian)													
<code>/var/log/pwdfail</code>													

To prevent a specific user to log in, their shell can be set either as:

<code>/bin/false</code>	user is forced to exit immediately
<code>/sbin/nologin</code>	user is prompted a message and forced to exit; the message is "This account is currently not available" or the contents of file <code>/etc/nologin.txt</code> if it exists

	Runlevel (SysV)	Target (Systemd)	Debian	Red Hat
		<code>halt.target</code>	System halt, no power off	
	0	<code>poweroff.target</code>	Shutdown	
	1	<code>rescue.target</code>	Single user / maintenance mode	
default runlevels	2		Multi-user mode (default)	Multi-user mode without network
	3	<code>multi-user.target</code>	Multi-user mode	Multi-user mode with network
	4		Multi-user mode	Unused, for custom use
	5	<code>graphical.target</code>	Multi-user mode	Multi-user mode with network and X (default)
	6	<code>reboot.target</code>	Reboot	
	S	<code>emergency.target</code>	Single user / maintenance mode with no mounted filesystems and no running services (usually accessed through runlevel 1)	
		<code>default.target</code>		System will always boot to this target; this is usually a symlink to <code>multi-user.target</code> or <code>graphical.target</code>

Systemd's target `runleveln.target` emulates a SysV's runlevel *n*.

```
runlevel
who -r
```

Display the previous and the current runlevel

```
init runlevel
telinit runlevel
```

Change to *runlevel*

```
systemctl get-default
systemctl set-default target
systemctl isolate target
systemctl emergency
systemctl rescue
systemctl -t target
```

Get the default target
Set *target* as the default target
Change to *target*
Change to maintenance single-user mode with only `/root` filesystem mounted
Change to maintenance single-user mode with only local filesystems mounted
List targets

To boot on the desired target a machine whose default target has become invalid (e.g. is symlinked to `reboot.target`), edit the GRUB 2 line and append `systemd.unit=desired.target` to the kernel parameters.

```
shutdown -h now
halt
poweroff
init 0
telinit 0
systemctl isolate shutdown.target
```

Shut down the system. Depending on the hardware, it will be either halted or powered off

```
shutdown -r now
reboot
init 6
telinit 6
systemctl isolate reboot.target
```

Reboot the system

```
shutdown
```

Shut down the system securely: all logged in users are notified via a message to their terminal, and login is disabled. Can only be run by the root user

```
shutdown -a
```

Non-root users that are listed in `/etc/shutdown.allow` can use this command to shut down the system

```
shutdown -h 16:00 message
```

Schedule a shutdown for 4 PM and send a warning message to all logged in users

```
shutdown -f
```

Skip `fsck` on reboot

```
shutdown -F
```

Force `fsck` on reboot

```
shutdown -c
```

Cancel a shutdown that has been already initiated

<code>/etc/init.d/service operation</code> <code>service service operation</code> <code>rcservice operation</code>	(RHEL 6 and earlier) (SUSE)	Perform the specified operation (<code>start</code> , <code>stop</code> , <code>status</code> , etc.) on the specified service
<code>update-rc.d service defaults</code> <code>chkconfig --add service</code>	(Debian) (RHEL 6 and earlier)	Add a service at boot
<code>update-rc.d -f service remove</code> <code>chkconfig --del service</code>	(Debian) (RHEL 6 and earlier)	Remove a service at boot
<code>update-rc.d -f service \</code> <code>start 30 2 3 4 5 . stop 70 0 1 6 .</code>		Add a service on the default runlevels; creates <code>s30</code> symlinks for starting the service and <code>k70</code> symlinks for stopping it
<code>chkconfig --levels 245 service on</code>		Add the service on runlevels 2 4 5
<code>chkconfig service on</code>		Add the service on default runlevels
<code>chkconfig service off</code>		Remove the service on default runlevels
<code>chkconfig service</code>		Check if the service is enabled on the current runlevel
<code>chkconfig service reset</code>		Reset the on/off state of the service for all runlevels to whatever the LSB specifies in the init script
<code>chkconfig service resetpriorities</code>		Reset the start/stop priorities of the service for all runlevels to whatever the LSB specifies in the init script
<code>chkconfig --list service</code>		Display current configuration of service (its status and the runlevels in which it is active)
<code>chkconfig</code> <code>chkconfig --list</code>		List all active services and their current configuration
<code>ls /etc/rcn.d</code> (Debian)		List services started on runlevel <i>n</i>

Supported service operations		
<code>start</code>	Start the service	Mandatory
<code>stop</code>	Stop the service	
<code>restart</code>	Restart the service (stop, then start)	
<code>status</code>	Display daemon PID and execution status	
<code>force-reload</code>	Reload configuration if service supports it, otherwise restart	
<code>condrestart</code> <code>try-restart</code>	Restart the service only if already running	Optional
<code>reload</code>	Reload the service configuration	

`systemctl operation service`

Perform the specified operation (*start, stop, status, etc.*) on the specified service (unit file)

`systemctl enable service`

Add the service on the current target

`systemctl disable service`

Remove the service on the current target

`systemctl is-enabled service`

Check if the service is enabled on the current target

`systemctl mask service`

Mask the service on the current target. This prevents the service to be enabled or started

`systemctl unmask service`

Unmask the service on the current target

`systemctl list-unit-files --type=service`

List all active services and their current configuration

`systemctl`

List loaded and active units

`systemctl --all`

List all units, including inactive ones

Systemd replaced SysV in RHEL 7 and later.

/etc/inittab	
<pre># The default runlevel. id:2:initdefault: # Boot-time system configuration/initialization script. # This is run first except when booting in emergency (-b) mode. si::sysinit:/etc/init.d/rcS # What to do in single-user mode. ~~:S:wait:/sbin/sulogin # /etc/init.d executes the S and K scripts upon change of runlevel. 10:0:wait:/etc/init.d/rc 0 11:1:wait:/etc/init.d/rc 1 12:2:wait:/etc/init.d/rc 2 13:3:wait:/etc/init.d/rc 3 14:4:wait:/etc/init.d/rc 4 15:5:wait:/etc/init.d/rc 5 16:6:wait:/etc/init.d/rc 6 # Normally not reached, but fall through in case of emergency. z6:6:respawn:/sbin/sulogin # /sbin/getty invocations for the runlevels. # Id field must be the same as the last characters of the device (after "tty"). 1:2345:respawn:/sbin/getty 38400 tty1 2:23:respawn:/sbin/getty 38400 tty2</pre>	

/etc/inittab describes which processes are started at bootup and during normal operation; it is read and executed by `init` at bootup.

All its entries have the form **id:runlevels:action:process**.

id	1-4 characters, identifies uniquely an entry. For gettys and other login processes it should be equal to the suffix of the corresponding tty	
runlevels	Runlevels for which the specified action must be performed. If empty, action is performed on all runlevels	
action	respawn wait once boot bootwait off ondemand initdefault sysinit powerfail powerwait powerfailnow powerokwait ctrlaltdel kbdrequest Process will be restarted when it terminates Process is started at the specified runlevel and <code>init</code> will wait for its termination (i.e. execution of further lines of <code>/etc/inittab</code> stops until the process exits) Process is executed once at the specified runlevel Process is executed at system boot. Runlevels field is ignored Process is executed at system boot and <code>init</code> will wait for its termination. Runlevels field is ignored Does nothing Process is executed when an on-demand runlevel (A, B, C) is called Specifies the default runlevel to boot on. Process field is ignored Process is executed at system boot, before any <code>boot</code> or <code>bootwait</code> entries. Runlevels field is ignored Process is executed when power goes down and a UPS activates. <code>init</code> will not wait for its termination Process is executed when power goes down and a UPS activates. <code>init</code> will wait for its termination Process is executed when power is down and the UPS battery is almost empty Process is executed when power has been restored from UPS Process is executed when <code>init</code> receives a SIGINT via <code>CTRL</code> <code>ALT</code> <code>DEL</code> Process is executed when a special key combination is pressed on console	
process	Process to execute. If prepended by a <code>+</code> , <code>utmp</code> and <code>wtmp</code> accounting will not be performed	

Directory	Content
/bin	Essential command binaries for all users
/boot	Bootloader files (OS loader, kernel images, initrd, etc.)
/dev	Virtual filesystem containing device nodes to devices and partitions
/etc	System configuration files and scripts
/home	Home directories for users
/lib	Libraries for the binaries in /bin and /sbin, kernel modules
/lost+found	Storage directory for recovered files in this partition
/media	Mount points for removable media
/mnt	Mount points for temporary filesystems
/net	Access to directory tree on different external NFS servers
/opt	Optional, large add-on application software packages
/proc	Virtual filesystem providing kernel and processes information
/root	Home directory for the root user
/run	Runtime variable data; replaces /var/run
/sbin	Essential system binaries, system administration commands
/srv	Data for services provided by the system
/sys	Virtual filesystem providing information about hotplug hardware devices
/tmp	Temporary files; deleted at reboot
/usr	User utilities and applications
/usr/bin	Non-essential command binaries for all users
/usr/include	C header files
/usr/lib	Libraries for the binaries in /usr/bin and /usr/sbin
/usr/local	Software installed locally
/usr/local/bin	Local software binaries
/usr/local/games	Local game binaries
/usr/local/include	Local C header files
/usr/local/lib	Local libraries for the binaries in /usr/local/bin and /usr/local/sbin
/usr/local/man	Local man pages
/usr/local/sbin	Local system binaries
/usr/local/share	Local architecture-independent hierarchy
/usr/local/src	Local source code
/usr/sbin	Non-essential system binaries (daemons and services)
/usr/share	Architecture-independent files (icons, fonts, documentation, etc.)
/usr/share/doc	Package-specific documentation not included in man pages
/usr/share/man	Man pages
/usr/share/info	Documentation in Info format
/usr/src	Source code for the current OS
/var	Variable files (logs, caches, mail spools, etc.)
/var/log	Logfiles
/var/opt	Variable files for the application software installed in /opt
/var/spool	Queued items to be processed (mail messages, cron jobs, print jobs, etc.)
/var/tmp	Temporary files that need to be stored for a longer time; preserved between reboots

The `manpage man hier` contains information about filesystem hierarchy.

The **superblock** contains information relative to the filesystem e.g. filesystem type, size, status, metadata structures.

The **Master Boot Record (MBR)** is a 512-byte program located in the first sector of the hard disk; it contains information about hard disk partitions and has the duty of loading the OS.

MBR has the following limits:

- max 4 primary partitions per hard disk, or 3 primary partitions + 1 extended partition; partitions numbered from 1 to 4
- max 11 logical partitions (inside the extended partition) per hard disk; partitions numbered from 5 to 15
- max disk size is 2 Tb

On recent systems, the MBR is replaced by the **GUID Partition Table (GPT)**. GPT does not differentiate between primary, extended, or logical partitions; furthermore, it practically has no limits on number and size of partitions.

When a partition is initialized, a **UUID (Universal Unique Identifier)**, which is a 128-bit hash number, is associated to it.

Almost all modern filesystems use **journaling**; in a journaling filesystem, the journal logs changes before committing them to the filesystem, which ensures faster recovery and less risk of corruption in case of a crash.

FUSE (Filesystem in Userspace) is an interface for userspace programs to export a filesystem to the Linux kernel, and is particularly useful for virtual file systems.

<code>fdisk /dev/sda</code>	Disk partitioning interactive tool
<code>fdisk -l /dev/sda</code>	List the partition table of device <code>/dev/sda</code>
<code>parted</code>	Disk partitioning interactive tool
<code>sfdisk /dev/sda</code>	Disk partitioning non-interactive tool
<code>cfdisk</code>	Disk partitioning tool with text-based UI
<code>gparted</code>	Disk partitioning tool with GUI
<code>gnome-disks</code>	Disk partitioning tool with GUI
<code>partprobe device</code>	Notify the kernel about disk and partition table changes. Otherwise, the changes will take place only after reboot
<code>hdparm -z device</code>	
<code>echo 1 > /sys/class/block/device/device/rescan</code>	

How to extend a virtual disk

1. Increase the size of an existing virtual disk
2. `parted -s /dev/sda "print"`

List the partitions of the virtual disk e.g. `/dev/sda`

```
Model: VMware Virtual disk (scsi)
Disk /dev/sda: 36.5GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
```

Number	Start	End	Size	File system	Name	Flags
1	1049kB	1075MB	1074MB	fat32	EFI System Partition	boot, esp
2	1075MB	2149MB	1074MB	xfs		
3	2149MB	36.5GB	34.4GB			lvm
3. `parted -s -a opt /dev/sda "resizepart 3 100%"`

Increase the last partition (in this case 3) so to take up all free space
4. `echo 1 > /sys/class/block/sda/device/rescan`

Notify the kernel about the partition increase
5. `lsblk`

Verify the partition increase

<code>mkfs -t fstype device</code>	Create a filesystem of the specified type on a partition (i.e. format the partition). <code>mkfs</code> is a wrapper utility for the actual filesystem-specific maker commands:
<code>mkfs.ext2</code>	aka <code>mke2fs</code>
<code>mkfs.ext3</code>	aka <code>mke3fs</code>
<code>mkfs.ext4</code>	
<code>mkfs.msdos</code>	aka <code>mkdosfs</code>
<code>mkfs.ntfs</code>	aka <code>mkntfs</code>
<code>mkfs.reiserfs</code>	aka <code>mkreiserfs</code>
<code>mkfs.jfs</code>	
<code>mkfs.xfs</code>	
<code>mkfs -t ext2 /dev/sda</code>	Create an ext2 filesystem on <code>/dev/sda</code>
<code>mkfs.ext2 /dev/sda</code>	
<code>mke2fs /dev/sda</code>	
<code>mke2fs -j /dev/sda</code>	Create an ext3 filesystem (ext2 with journaling) on <code>/dev/sda</code>
<code>mkfs.ext3 /dev/sda</code>	
<code>mke3fs /dev/sda</code>	
<code>mkfs -t msdos /dev/sda</code>	Create a MS-DOS filesystem on <code>/dev/sda</code>
<code>mkfs.msdos /dev/sda</code>	
<code>mkdosfs /dev/sda</code>	
<code>blockdev --getbsz /dev/sda1</code>	Get the block size of the specified partition
<code>wipefs device</code>	List all visible filesystems and their signatures' offsets
<code>wipefs -a device</code>	Erase filesystem or RAID signatures (magic strings i.e. metadata) from the device to make the filesystem invisible from <code>blkid</code>
<code>blkdiscard options device</code>	Discard device sectors, wiping the data they contain. Useful for SSDs
<code>resize2fs options device size</code>	Resize an ext2/ext3/ext4 filesystem
<code>file -s /dev/sda</code>	Show information about device <code>/dev/sda</code> e.g. whether it uses MBR
<code>fsfreeze -f mountpoint</code>	Freeze (suspend) access to a mounted filesystem, creating a stable image on disk. Used with hardware RAID that support snapshots. Unnecessary with LVM devices
<code>fsfreeze -u mountpoint</code>	Unfreeze (resume) access to a mounted filesystem
<code>blkid /dev/sda1</code>	Print the UUID of the specified partition
<code>blkid -L /boot</code>	Print the UUID of the specified partition, given its label
<code>blkid -U 652b786e-b87f-49d2-af23-8087ced0c667</code>	Print the name of the specified partition, given its UUID
<code>findfs UUID=652b786e-b87f-49d2-af23-8087ced0c667</code>	Print the name of the specified partition, given its UUID
<code>findfs LABEL=/boot</code>	Print the name of the specified partition, given its label
<code>e2label /dev/sda1</code>	Print the label of the specified partition

<code>fsck device</code>	Check and repair a Linux filesystem. The filesystem must be unmounted; running fsck on a mounted filesystem, even if mounted read-only, risks damaging it. Corrupted files will be placed into the <code>/lost+found</code> directory of the partition. fsck is a wrapper utility for the actual filesystem-specific checker commands: <pre>fsck.ext2 aka e2fsck fsck.ext3 aka e2fsck fsck.ext4 aka e2fsck fsck.msdos fsck.vfat fsck.cramfs fsck.minix</pre>
<code>fsck</code> <code>fsck -As</code>	Check and repair serially all filesystems listed in <code>/etc/fstab</code>
<code>fsck -f /dev/sda1</code>	Force a filesystem check on <code>/dev/sda1</code> even if fsck thinks it is not necessary
<code>fsck -y /dev/sda1</code>	During filesystem repair, do not ask questions and assume that the answer is always yes
<code>fsck -n /dev/sda1</code>	Perform safely a filesystem check on a mounted filesystem, only reporting errors without trying to correct them. Not recommended as the results will not be accurate, and some types of filesystem do not even support this option
<code>fsck.ext2 -c /dev/sda1</code>	Check an ext2 filesystem, running the <code>badblocks</code> command to mark all bad blocks and add them to the bad block inode, so that they will not be allocated to files or directories
<code>touch /forcefsck</code>	Force a filesystem check after next reboot (RHEL 6 and earlier)

mount	Display the currently mounted filesystems, and their mount options.
cat /proc/mounts	The commands <code>mount</code> and <code>umount</code> maintain in <code>/etc/mtab</code> a database of currently mounted filesystems, but <code>/proc/mounts</code> is authoritative
cat /etc/mtab	
mount -a	Mount all devices listed in <code>/etc/fstab</code> , except those indicated as <code>noauto</code>
mount -t ext3 /dev/sda /mnt	Mount a Linux-formatted disk. The mount point (directory) must exist
mount -t msdos /dev/fd0 /mnt	Mount a MS-DOS filesystem floppy disk to mount point <code>/mnt</code>
mount /dev/fd0	Mount a floppy disk. <code>/etc/fstab</code> must contain an entry for <code>/dev/fd0</code>
mount -o remount,rw /	Remount the root directory as read-write, supposing it was mounted read-only. Useful to change flags (in this case, read-only to read-write) for a mounted filesystem that cannot be unmounted at the moment
mount -o nolock 10.7.7.7:/export/ /mnt/nfs	Mount a NFS share without running NFS daemons. Useful during system recovery
mount -t iso9660 -o ro,loop=/dev/loop0 cd.img /mnt/cdrom	Mount a CD-ROM ISO9660 image file like a CD-ROM (via the loop device)
umount /dev/fd0	Unmount a floppy disk that was mounted on <code>/mnt</code> (device must not be busy)
umount /mnt	
umount -l /dev/fd0	Unmount the floppy disk as soon as it is not in use anymore
mountpoint /mnt	Tell if a directory is a mount point
findmnt	List all mounted filesystems
findmnt /dev/sda	Find the filesystem, given a specified device or mount point
findmnt /mnt	
eject /dev/fd0	Eject a removable media device
eject /mnt	

Partition types					
0x00	Empty	0x4e	QNX4.x 2 nd part	0xa8	Darwin UFS
0x01	FAT12	0x4f	QNX4.x 3 rd part	0xa9	NetBSD
0x02	XENIX root	0x50	OnTrack DM	0xab	Darwin boot
0x03	XENIX usr	0x51	OnTrack DM6 Aux1	0xaf	HFS / HFS+
0x04	FAT16 < 32Mb	0x52	CP/M	0xb7	BSDI fs
0x05	Extended	0x53	OnTrack DM6 Aux3	0xb8	BSDI swap
0x06	FAT16	0x54	OnTrackDM6	0xbb	Boot Wizard hidden
0x07	HPFS / NTFS / exFAT	0x55	EZ-Drive	0xbe	Solaris boot
0x08	AIX	0x56	Golden Bow	0xbf	Solaris
0x09	AIX bootable	0x5c	Priam Edisk	0xc1	DRDOS/sec (FAT-12)
0x0a	OS/2 Boot Manager	0x61	SpeedStor	0xc4	DRDOS/sec (FAT-16 < 32Mb)
0x0b	W95 FAT32	0x63	GNU HURD or SysV	0xc6	DRDOS/sec (FAT-16)
0x0c	W95 FAT32 (LBA)	0x64	Novell Netware 286	0xc7	Syrinx
0x0e	W95 FAT16 (LBA)	0x65	Novell Netware 386	0xda	Non-FS data
0x0f	W95 extended (LBA)	0x70	DiskSecure Multi-Boot	0xdb	CP/M, CTOS, etc.
0x10	OPUS	0x75	PC/IX	0xde	Dell Utility
0x11	Hidden FAT12	0x80	Old Minix	0xdf	BootIt
0x12	Compaq diagnostics	0x81	Minix / old Linux	0xe1	DOS access
0x14	Hidden FAT16 < 32Mb	0x82	Linux swap / Solaris	0xe3	DOS R/O
0x16	Hidden FAT16	0x83	Linux	0xe4	SpeedStor
0x17	Hidden HPFS/NTFS	0x84	OS/2 hidden C: drive	0xeb	BeOS fs
0x18	AST SmartSleep	0x85	Linux extended	0xee	GPT
0x1b	Hidden W95 FAT32	0x86	NTFS volume set	0xef	EFI (FAT-12/16/32)
0x1c	Hidden W95 FAT32 (LBA)	0x87	NTFS volume set	0xf0	Linux/PA-RISC boot
0x1e	Hidden W95 FAT16 (LBA)	0x88	Linux plaintext	0xf1	SpeedStor
0x24	NEC DOS	0x8e	Linux LVM	0xf4	SpeedStor
0x27	Hidden NTFS WinRE	0x93	Amoeba	0xf2	DOS secondary
0x39	Plan 9	0x94	Amoeba BBT	0xfb	VMware VMFS
0x3c	PartitionMagic recovery	0x9f	BSD/OS	0xfc	VMware VMKCORE
0x40	Venix 80286	0xa0	IBM Thinkpad hibernation	0xfd	Linux RAID autodetect
0x41	PPC PReP Boot	0xa5	FreeBSD	0xfe	LANstep
0x42	SFS	0xa6	OpenBSD	0xff	BBT
0x4d	QNX4.x	0xa7	NeXTSTEP		

The command `sfdisk -T` prints the above list of partition IDs and names.

Most common Linux-supported filesystems	
ext2	The oldest Linux ext filesystem, without journaling
ext3	ext2 with journaling
ext4	Linux journaling filesystem, an upgrade from ext3
Reiserfs	Journaling filesystem
XFS	Journaling filesystem, developed by SGI. Offers more performance and scalability than ext4
JFS	Journaling filesystem, developed by IBM
Btrfs	B-tree filesystem, developed by Oracle
msdos	DOS filesystem, supporting only 8-char filenames
umsdos	Extended DOS filesystem used by Linux, compatible with DOS
fat32	MS Windows FAT filesystem
vfat	Extended DOS filesystem, with support for long filenames
ntfs	Replacement for fat32 and vfat filesystems
minix	Native filesystem of the MINIX OS, created for didactic purposes
iso9660	CD-ROM filesystem
cramfs	Compressed RAM disk
nfs	Network filesystem, used to access files on remote machines
SMB	Server Message Block, used to mount MS Windows network shares
proc	Pseudo filesystem, used as an interface to the kernel data structures
swap	Pseudo filesystem, Linux swap area

The **swap** space is an area on disk (a file or a partition) used as a RAM extension. When there is not enough free physical RAM for a process, inactive pages in memory are temporarily **swapped out** of memory to disk, to later be **swapped in** to memory when RAM resources are available again. If both RAM and swap space become nearly full, the system may get clogged by spending all the time paging blocks of memory back and forth between RAM and swap (**thrashing**). The amount of RAM plus the swap is defined as the **virtual memory**.

In Linux, a swap partition is usually preferred over a swap file. While a swap file can be resized more easily, it cannot be used for hibernation; this because the system must first locate the swap file's header, but in order to do so the filesystem containing the swap file must be mounted, and journaled filesystems such as ext3 or ext4 cannot be mounted during resume from disk. Also, in older Linux versions a swap partition used to have faster disk access and less fragmentation than a swap file; in modern Linux versions the difference is negligible.

The swap partition is listed as filesystem type 0x82; however, it is not a filesystem, but a raw addressable memory space with no structure. For this reason it does not appear in the output of `mount` or `df` commands.

A swap partition can be created via any partitioning tool e.g. `fdisk`.

<code>dd if=/dev/zero of=/swapfile bs=1024 count=512000</code>	Create a 512-Mb swap file
<code>mkswap /swapfile</code>	Initialize a (already created) swap file or partition
<code>swapon /swapfile</code>	Enable a swap file or partition, thus telling the kernel that it can use it now
<code>swapoff /swapfile</code>	Disable a swap file or partition
<code>swapon -s</code> <code>cat /proc/swaps</code> <code>cat /proc/meminfo</code> <code>free</code> <code>top</code>	Show the sizes of total and used swap areas

How to extend a LVM swap partition

1. <code>lvs</code>	Determine the name of the swap Logical Volume
2. <code>swapoff /dev/volgroup0/swap_lv</code>	Turn off the swap volume
3. <code>lvresize -L+1G /dev/volgroup0/swap_lv</code>	Extend the swap volume with an additional 1 Gb of space
4. <code>mkswap /dev/volgroup0/swap_lv</code>	Format the swap volume
5. <code>swapon /dev/volgroup0/swap_lv</code>	Turn on the swap volume

/etc/fstab					
# <filesystem>	<mount point>	<type>	<options>	<dump>	<pass>
/dev/sda2	/	ext2	defaults	0	1
/dev/sdb1	/home	ext2	defaults	1	2
/dev/cdrom	/media/cdrom	auto	ro,noauto,user,exec	0	0
/dev/fd0	/media/floppy	auto	rw,noauto,user,sync	0	0
proc	/proc	proc	defaults	0	0
/dev/hda1	swap	swap	pri=42	0	0
nfsserver:/dirs	/mnt	nfs	intr	0	0
//smbserver/jdoe	/shares/jdoe	cifs	auto,credentials=/etc/smbcreds	0	0
LABEL=/boot	/boot	ext2	defaults	0	0
UUID=652b786e-b87f-49d2-af23-8087ced0c667	/test	ext4	errors=remount-ro,noatime	0	0

/etc/fstab contains information about filesystems, including all filesystems that must be automatically mounted at bootup.

filesystem	Device or partition. The filesystem can be identified either by its name, label, or UUID	
mount point	Directory on which the partition will be mounted	
type	Filesystem type, or <code>auto</code> if detected automatically	
options	defaults	Use the default options. The default options depend on the filesystem type and can be found via the command: <code>tune2fs -l device grep "Default mount options"</code> Most common default options: <code>rw, suid, dev, auto, nouser, exec, async</code>
	ro	Mount read-only
	rw	Mount read-write (default)
	suid	Permit SUID and SGID bit operations (default)
	nosuid	Do not permit SUID and SGID bit operations
	dev	Interpret block special devices on the filesystem (default)
	nodev	Do not interpret block special devices on the filesystem
	auto	Mount automatically at bootup, or when command <code>mount -a</code> is given (default)
	noauto	Mount only if explicitly demanded
	user	Partition can be mounted by any user
	nouser	Partition can be mounted only by the root user (default)
	exec	Binaries contained on the partition can be executed (default)
	noexec	Binaries contained on the partition cannot be executed
	sync	Write files immediately to the partition
	async	Buffer write operations and commit them at once later, or when device is unmounted (default)
	noatime	Do not update atime (access time) information for files. This results in a performance improvement because the system does not need anymore to do filesystem writes for files which are just being read
	nodiratime	Do not update atime (access time) information for directories
	acl	Support ACLs on files contained in the partition
	context="context"	Apply a specific SELinux context to the mount
Other specific options apply to specific partition types (e.g. NFS or Samba)		
dump	Options for the <code>dump</code> backup utility. 0 = do not backup	
pass	Order in which the filesystem must be checked by <code>fsck</code> . 0 = do not check. The filesystem root should be 1	

<code>df</code>	Report filesystem disk space usage
<code>df -h</code>	Report filesystem disk space usage in human-readable output
<code>df <i>directory</i></code>	Shows on which device the specified <i>directory</i> is mounted
<code>du <i>directory</i></code>	Report disk usage, as the size of each file contained in <i>directory</i> , in Kb
<code>du -s <i>directory</i></code>	Show the total sum of the sizes of all files contained in <i>directory</i>
<code>du -h <i>directory</i></code>	Report disk usage in human-readable output
<code>du -hs * sort -hr</code>	Print out all files and directories in the current directory, ordered by size (the largest first), in human-readable output
<code>du -a /<i>path</i> sort -nr head</code>	Print out the 10 biggest files and directories under <i>path</i>
<code>find /<i>path</i> -type f -exec du -Sh {} + \</code> <code> sort -hr head</code>	Print out the 10 biggest files under <i>path</i>
<code>ncdu</code>	Disk usage analyzer with ncurses UI
<code>duf</code>	Disk usage analyzer that shows the results in a table format
<code>lsblk</code>	List information about all available block devices
<code>lsscsi</code>	List information about all SCSI devices
<code>sync</code>	Flush the buffer and commit all pending writes. To improve performance of Linux filesystems, many write operations are buffered in RAM and written at once; writes are done in any case before unmount, reboot, or shutdown
<code>chroot /<i>path/to/newrootdir</i> <i>command</i></code>	Run a command in a chroot jail (i.e. in a new root directory). The command process will be unable to access files outside the chroot jail
<code>chroot /mnt/sysimage</code>	Start a shell with <code>/mnt/sysimage</code> as filesystem root. Useful during system recovery when the machine has been booted from a removable media; this device is defined as the filesystem root and often needs to be changed to perform operations on the machine
<code>mknod <i>device type major minor</i></code>	Create a character or block special file, allocating the proper inode. Useful if experiencing filesystem problems during system recovery
<code>losetup <i>options</i></code>	Associate/disassociate a loop device with a regular file or block device, or query the status of a loop device
<code>multipath <i>options device</i></code>	Detect and aggregate multiple I/O paths (SAN connections) to a device

<code>tune2fs options device</code>	Adjust tunable filesystem parameters on ext2/ext3/ext4 filesystems
<code>tune2fs -l /dev/sda1</code>	List the contents of the filesystem superblock
<code>tune2fs -j /dev/sda1</code>	Add a journal to this ext2 filesystem, making it an ext3
<code>tune2fs -m 1 /dev/sda1</code>	Reserve 1% of the partition size to privileged processes. This space (5% by default) is reserved to avoid filesystem fragmentation and to allow privileged processes to continue to run correctly even when the partition is full
<code>tune2fs -C 7 /dev/sda1</code>	Set the mount count of the filesystem to 7
<code>tune2fs -c 20 /dev/sda1</code>	Set the filesystem to be checked by fsck after 20 mounts
<code>tune2fs -i 15d /dev/sda1</code>	Set the filesystem to be checked by fsck each 15 days
<code>tune2fs -l /dev/sda1 \</code> <code>grep "Default mount options"</code>	Print the default mount options for the filesystem

Both mount-count-dependent and time-dependent checking are enabled by default for all hard drives on Linux, to avoid the risk of filesystem corruption going unnoticed.

<code>dumpe2fs options device</code>	Dump ext2/ext3/ext4 filesystem information
<code>dumpe2fs -h /dev/sda1</code>	Display filesystem's superblock information (number of mounts, last checks, UUID, etc.)
<code>dumpe2fs /dev/sda1 grep -i superblock</code>	Display locations of superblock (primary and backup) of filesystem
<code>dumpe2fs -b /dev/sda1</code>	Display blocks that are marked as bad in the filesystem
<code>debugfs device</code>	Interactive ext2/ext3/ext4 filesystem debugger
<code>debugfs -w /dev/sda1</code>	Debug <code>/dev/sda1</code> in read-write mode (by default, <code>debugfs</code> accesses the device in read-only mode)
<code>e2freefrag /dev/sda1</code>	Report free space fragmentation on an ext2/ext3/ext4 filesystem
<code>filefrag file</code>	Display the number of extents into which a file is fragmented
<code>e4defrag -c directory</code>	Report filesystem fragmentation
<code>hdparm</code>	Get or set drive parameters for SATA/IDE devices
<code>hdparm -g /dev/hda</code>	Display drive geometry (cylinders, heads, sectors) of <code>/dev/hda</code>
<code>hdparm -i /dev/hda</code>	Display identification information for <code>/dev/hda</code>
<code>hdparm -tT /dev/hda</code>	Perform disk read benchmarks on the <code>/dev/hda</code> drive
<code>hdparm -p 12 /dev/hda</code>	Reprogram IDE interface chipset of <code>/dev/hda</code> to mode 4. Warning: using an unsupported mode can cause filesystem corruption
<code>sdparm</code>	Access drive parameters for SCSI devices

Many hard drives feature the **Self-Monitoring, Analysis and Reporting Technology (SMART)** whose purpose is to monitor the reliability of the drive, predict drive failures, and carry out different types of drive self-tests. The `smartd` daemon attempts to poll this information from all drives every 30 minutes, logging all data to `syslog`.

<code>smartctl -a /dev/sda</code>	Print SMART information for drive <code>/dev/sda</code>
<code>smartctl -s off /dev/sda</code>	Disable SMART monitoring and log collection for drive <code>/dev/sda</code>
<code>smartctl -t long /dev/sda</code>	Begin an extended SMART self-test on drive <code>/dev/sda</code>

inotify (inode notify) is a Linux kernel subsystem that monitors changes to the filesystem and notifies applications about these changes; it can hence be used to reload configuration files, check log changes, synchronize files, etc. It replaced **dnotify** in kernel version 2.6.13.

`inotifywatch` Listen for filesystem events, and output a summary of the events received on each file or directory

`inotifywait` Wait for changes to files

incron (inotify cron) is a system similar to cron, except that it handles filesystem events rather than time period events.

`incrontab` Manipulate incron tables

<code>xfs_growfs options mountpoint</code>	Expand an XFS filesystem. Filesystem shrinking is not supported by XFS
<code>xfs_info /dev/sda1</code> <code>xfs_growfs -n /dev/sda1</code>	Print XFS filesystem geometry
<code>xfs_check options device</code>	Check XFS filesystem consistency
<code>xfs_repair options device</code>	Repair a damaged or corrupt XFS filesystem
<code>xfs_db -c frag -r device</code>	Display the level of fragmentation of a XFS filesystem
<code>xfs_fsr device</code>	Defragment a XFS filesystem
<code>xfsdump -v silent -f /dev/tape /</code>	Dump the root of a XFS filesystem to tape, with the lowest verbosity. Incremental and resumed dumps are stored in the inventory database <code>/var/lib/xfsdump/inventory</code>
<code>xfsdump -J - / xfsrestore -J - /new</code>	Copy the contents of a XFS filesystem to another directory, without updating the inventory database
<code>xfsrestore -f /dev/tape /</code>	Restore a XFS filesystem from tape
<code>reiserfstune options device</code>	Adjust tunable filesystem parameters on ReiserFS filesystem
<code>debugreiserfs device</code>	Interactive ReiserFS filesystem debugger
<code>mkisofs -r -o cdrom.img data/</code>	Create a CD-ROM image with a ISO9660 filesystem from the contents of the target directory. This command enables Rock Ridge extension, which contains the original file information for MS Windows 8.3 filenames e.g. permissions and filename; it also sets all content on CD to be publicly readable, instead of inheriting the permissions from the original files. Other CD-ROM filesystems are UDF (Universal Disk Format) and HFS (Hierarchical File System). Other CD-ROM filesystem extensions are MS Joliet, that allows creating CD-ROMs that are more MS Windows compatible, and El Torito, to create bootable CD-ROMs
<code>mkudffs /dev/hda</code>	Create a UDF filesystem
<code>udffsck /dev/hda</code>	Check a UDF filesystem
<code>wrudf /dev/hda</code>	Maintain a UDF filesystem. Provides an interactive shell
<code>cdrwtool -d /dev/sr0 operation</code>	Manage a CD-RW drive (e.g. blank the disc, set read/write speed, set write parameters, display disc track info)

AutoFS is a client-side service that allows automounting of filesystems, even for nonprivileged users. AutoFS is composed of the `autofs` kernel module that monitors specific directories for attempts to access them; in this case, the kernel module signals the `automount` userspace daemon, which mounts the directory when it needs to be accessed and unmounts it when is no longer accessed.

Mounts managed by AutoFS should not be mounted/unmounted manually or via `/etc/fstab`, to avoid inconsistencies.

AutoFS configuration files	
<code>/etc/sysconfig/autofs</code>	AutoFS configuration file.
<code>/etc/auto.master</code>	Master map file for AutoFS. Each line is an indirect map, and each map file stores the configuration for the automounting of the subdirectory. The <code>-hosts</code> map tells AutoFS to mount/unmount automatically any export from the NFS server <code>nfsserver</code> when the directory <code>/net/nfsserver/</code> is accessed. <pre># mount point map options /net -hosts /- /etc/auto.direct /misc /etc/auto.misc /home /etc/auto.home --timeout=60</pre>

AutoFS map files	
<code>/etc/auto.direct</code>	Direct map file for automounting of a NFS share. <pre># dir filesystem /mydir nfsserver1.foo.org:/myshare</pre>
<code>/etc/auto.misc</code>	Indirect map file for automounting of directory <code>/misc</code>. <pre># subdir options filesystem public -ro,soft,intr ftp.example.org:/pub cd -fstype=iso9660,ro,nosuid,nodev :/dev/cdrom</pre>
<code>/etc/auto.home</code>	Indirect map file for automounting of directory <code>/home</code> on a NFS share. The <code>*</code> wildcard matches any subdirectory the system attempts to access, and the <code>&</code> variable takes the value of the match. <pre># subdir options filesystem * -rw,soft,intr nfsserver2.bar.org:/home/&</pre>

RAID levels		
Level	Description	Storage capacity
RAID 0	Striping (data is written across all member disks). High I/O but no redundancy	Sum of the capacity of member disks
RAID 1	Mirroring (data is mirrored on all disks). High redundancy but high cost	Capacity of the smaller member disk
RAID 4	Parity (for fault tolerance) on a single disk. I/O bottleneck unless coupled to write-back caching	Sum of the capacity of member disks, minus one
RAID 5	Parity distributed across all disks. Can sustain one disk crash	Sum of the capacity of member disks, minus one
RAID 6	Double parity distributed across all disks. Can sustain two disk crashes	Sum of the capacity of member disks, minus two
RAID 10 (1+0)	Striping + mirroring. High redundancy but high cost	Capacity of the smaller member disk
Linear RAID	Data written sequentially across all disks. No redundancy	Sum of the capacity of member disks

```
mdadm -C /dev/md0 -l 5 \
-n 3 /dev/sdb1 /dev/sdc1 /dev/sdd1 \
-x 1 /dev/sde1
```

Create a RAID 5 array from three partitions and a spare.
Partitions type must be set to 0xFD.
Once the RAID device has been created, it must be formatted e.g. via
`mke2fs -j /dev/md0`

```
mdadm --manage /dev/md0 -f /dev/sdd1
mdadm --manage /dev/md0 -r /dev/sdd1
```

Mark a drive as faulty, before removing it
Remove a drive from the RAID array.
The faulty drive can then be physically removed

```
mdadm --manage /dev/md0 -a /dev/sdd1
```

Add a drive to the RAID array.
To be run after the faulty drive has been physically replaced

```
mdadm --misc -Q /dev/sdd1
mdadm --misc -D /dev/md0
mdadm --misc -o /dev/md0
mdadm --misc -w /dev/md0
```

Display information about a device
Display detailed information about the RAID array
Mark the RAID array as read-only
Mark the RAID array as read & write

```
cat /proc/mdstat
```

Display information about RAID arrays and devices

```
/etc/mdadm.conf
```

Configuration file for the `mdadm` command

```
DEVICE /dev/sdb1 /dev/sdc1 /dev/sdd1 /dev/sde1
ARRAY /dev/md0 level=raid5 num-devices=3
        UUID=0098af43:812203fa:e665b421:002f5e42
        devices=/dev/sdb1,/dev/sdc1,/dev/sdd1,/dev/sde1
```

The **bootloader** (aka **boot manager** or **bootstrap loader**) is the program responsible for booting a machine. The standard bootloader on Linux distributions is **GRUB (Grand Unified Bootloader)**. The latest version is **GRUB 2**; the previous version is known as **GRUB Legacy** or **GRUB 1**.

Non-GRUB bootloaders										
LILO (Linux Loader)		Obsolete. Small bootloader that can be placed in the MBR or the boot sector of a partition. The configuration file is <code>/etc/lilo.conf</code> (run <code>/sbin/lilo</code> afterwards to validate changes). The version for EFI hardware is ELILO.								
Syslinux	SYSLINUX	Able to boot from FAT and NTFS filesystems e.g. floppy disks and USB drives. Used for boot floppy disks, rescue floppy disks, and Live USBs.								
	ISOLINUX	Able to boot from CD-ROM ISO 9660 filesystems. Used for Live CDs and bootable install CDs. The CD must contain the following files: <table><tr><td><code>isolinux/isolinux.bin</code></td><td>ISOLINUX image, from the SYSLINUX distro</td></tr><tr><td><code>boot/isolinux/isolinux.cfg</code></td><td>ISOLINUX configuration</td></tr><tr><td><code>images/</code></td><td>Floppy images to boot</td></tr><tr><td><code>kernel/memdisk</code></td><td></td></tr></table> and can be burnt with the command: <pre>mkisofs -o output.iso -b isolinux/isolinux.bin -c isolinux/boot.cat \ -no-emul-boot -boot-load-size 4 -boot-info-table cd_root_dir</pre>	<code>isolinux/isolinux.bin</code>	ISOLINUX image, from the SYSLINUX distro	<code>boot/isolinux/isolinux.cfg</code>	ISOLINUX configuration	<code>images/</code>	Floppy images to boot	<code>kernel/memdisk</code>	
	<code>isolinux/isolinux.bin</code>	ISOLINUX image, from the SYSLINUX distro								
	<code>boot/isolinux/isolinux.cfg</code>	ISOLINUX configuration								
<code>images/</code>	Floppy images to boot									
<code>kernel/memdisk</code>										
PXELINUX	Able to boot from PXE (Pre-boot eXecution Environment). PXE uses DHCP or BOOTP to enable basic networking, then uses TFTP to download a bootstrap program that loads and configures the kernel. Used for Linux installations from a central server or network boot of diskless workstations. The boot TFTP server must contain the following files: <table><tr><td><code>/tftpboot/pxelinux.0</code></td><td>PXELINUX image, from the SYSLINUX distribution</td></tr><tr><td><code>/tftpboot/pxelinux.cfg/</code></td><td>Directory containing a configuration file for each machine. A machine with Ethernet MAC address 88:99:AA:BB:CC:DD and IP address 192.0.2.91 (C000025B in hexadecimal) will search for its configuration filename in this order: 01-88-99-aa-bb-cc-dd C000025B C000025 C00002 C0000 C000 C00 C0 C default</td></tr></table>	<code>/tftpboot/pxelinux.0</code>	PXELINUX image, from the SYSLINUX distribution	<code>/tftpboot/pxelinux.cfg/</code>	Directory containing a configuration file for each machine. A machine with Ethernet MAC address 88:99:AA:BB:CC:DD and IP address 192.0.2.91 (C000025B in hexadecimal) will search for its configuration filename in this order: 01-88-99-aa-bb-cc-dd C000025B C000025 C00002 C0000 C000 C00 C0 C default					
<code>/tftpboot/pxelinux.0</code>	PXELINUX image, from the SYSLINUX distribution									
<code>/tftpboot/pxelinux.cfg/</code>	Directory containing a configuration file for each machine. A machine with Ethernet MAC address 88:99:AA:BB:CC:DD and IP address 192.0.2.91 (C000025B in hexadecimal) will search for its configuration filename in this order: 01-88-99-aa-bb-cc-dd C000025B C000025 C00002 C0000 C000 C00 C0 C default									
EXTLINUX	General-purpose bootloader like LILO or GRUB. Now merged with SYSLINUX.									

In GRUB 2, the GRUB bootstrap code i.e. GRUB Stage 1 (446 bytes) is stored in the 512-byte MBR; the MBR also contains the partition table (64 bytes) and the boot signature (2 bytes).

Once booted, GRUB Stage 1 locates and executes GRUB Stage 1.5.

GRUB Stage 1.5 contains common filesystem drivers necessary to locate GRUB Stage 2.

GRUB Stage 2 accesses the GRUB 2 configuration and commands stored in `/boot/grub2`. It loads one of the available Linux kernels into RAM and passes control to it.

`/boot/grub/grub.cfg` or `/boot/grub2/grub.cfg` GRUB 2 configuration file

```
# Linux Red Hat
menuentry "Fedora 2.6.32" {
    set root=(hd0,1)
    linux /vmlinuz-2.6.32 ro root=/dev/hda5 mem=2048M
    initrd /initrd-2.6.32
}

# Linux Debian
menuentry "Debian 2.6.36-experimental" {
    set root=(hd0,1)
    linux (hd0,1)/bzImage-2.6.36-experimental ro root=/dev/hda6
}

# MS Windows
menuentry "MS Windows" {
    set root=(hd0,2)
    chainloader +1
}
```

The GRUB 2 configuration file must not be edited manually. Instead, one must edit the files in `/etc/grub.d/` (these are scripts that will be run in order) and the file `/etc/default/grub` (configuration file for menu display settings), then run `update-grub` (Debian) or `grub2-mkconfig` (Red Hat) which will recreate this configuration file.

GRUB 2 configuration - Common kernel parameters

<code>root=</code>	Specify the location of the filesystem root. This is a required parameter
<code>ro</code>	Mount read-only on boot
<code>rw</code>	Mount read-write on boot
<code>quiet</code>	Disable non-critical kernel messages during boot
<code>debug</code>	Enable kernel debugging
<code>rhgb</code>	Red Hat graphical boot. GUI mode booting screen
<code>splash</code>	GUI mode booting screen
<code>single</code>	Boot in single-user mode (runlevel 1)
<code>emergency</code>	Emergency mode: after the kernel is booted, run <code>sulogin</code> (single-user login) which asks for the root password for system maintenance, then run a Bash shell. Does not load <code>init</code> or any daemon or configuration setting
<code>init=/bin/bash</code>	Run a Bash shell (may also be any other executable) instead of <code>init</code>
<code>rd.break</code>	Interrupt the boot process and land to a shell in ramdisk environment
<code>fsck.mode=</code>	Filesystem check at boot. Possible options: <code>auto</code> Filesystem check is done only if necessary. Default <code>force</code> Force a full filesystem check <code>skip</code> Skip any filesystem check
<code>fsck.repair=</code>	Filesystem check at boot. Possible options: <code>preen</code> Automatically repair filesystem problems that can be safely fixed <code>yes</code> Answer yes to all questions <code>no</code> Answer no to all questions
<code>rd.lvm.lv=</code>	Only activate the specified LVM volumes. Can be specified multiple times

The GRUB menu, presented at startup, allows to choose the OS or kernel to boot:

ENTER	Boot the currently selected GRUB entry
C	Get a GRUB command line
E	Edit the selected GRUB entry (e.g. to edit kernel parameters in order to boot in single-user emergency mode, or to change IRQ or I/O port of a device driver compiled in the kernel)
B	Boot the currently selected GRUB entry. This is usually done after finishing modifying the entry
P	Bring up the GRUB password prompt. Necessary if a GRUB password has been set

<code>grub2-mkconfig</code>	Generate a GRUB configuration file
<code>grub2-mkconfig -o /boot/grub2/grub.cfg</code>	Regenerate the GRUB configuration file for a BIOS system
<code>grub2-mkconfig -o /boot/efi/EFI/centos/grub.cfg</code>	Regenerate the GRUB configuration file for an UEFI system
<code>update-grub</code>	Generate a GRUB configuration file
<code>grub-install /dev/sda</code>	Install GRUB on the first SATA drive
<code>grub</code>	Access the GRUB shell
<code>grub2-set-default 1</code>	Set GRUB to automatically boot the second entry in the GRUB menu
<code>grub2-editenv list</code>	Display the current GRUB menu entry that is automatically booted
<code>grub2-editenv - unset menu_auto_hide</code>	Set the GRUB to always show up at boot
<code>grub-emu</code>	GUI tool that allows to see the current GRUB menu without rebooting
<code>grubby</code>	Tool for editing the configuration files for GRUB, LILO, and ELILO

<code>/boot/grub/device.map</code>	If present, this file maps Linux device filenames to BIOS drives: (fd0) /dev/fd0 (hd0) /dev/hda
------------------------------------	---

/boot/grub/menu.lst or /boot/grub/grub.conf GRUB Legacy configuration file

```
timeout 10    # Boot the default kernel after 10 seconds
default 0     # Default kernel is 0

# Section 0: Linux boot
title    Debian    # Menu item to show on GRUB bootmenu
root     (hd0,0)   # root filesystem is /dev/hda1
kernel   /boot/vmlinuz-2.6.24-19-generic root=/dev/hda1 ro quiet splash
initrd   /boot/initrd.img-2.6.24-19-generic

# Section 1: MS Windows boot
title    Microsoft Windows XP
root     (hd0,1)   # root filesystem is /dev/hda2
savedefault
makeactive    # set the active flag on this partition
chainloader +1    # read 1 sector from start of partition and run

# Section 2: Firmware/BIOS update from floppy disk
title    Firmware update
kernel   /memdisk    # boot a floppy disk image
initrd   /floppy-img-7.7.7
```

GRUB Legacy shell commands	
<code>blocklist file</code>	Print the block list notation of a file
<code>boot</code>	Boot the loaded OS
<code>cat file</code>	Show the contents of a file
<code>chainloader file</code>	Chainload another bootloader
<code>cmp file1 file2</code>	Compare two files
<code>configfile file</code>	Load a configuration file
<code>debug</code>	Toggle debugging mode
<code>displayapm</code>	Display APM BIOS information
<code>displaymem</code>	Display memory configuration
<code>embed stage device</code>	Embed Stage 1.5 in the device
<code>find file</code>	Find a file
<code>fstest</code>	Toggle filesystem test mode
<code>geometry drive</code>	Print information on a drive geometry
<code>halt</code>	Shut down the system
<code>help command</code>	Show help for a command, or the available commands
<code>imsprobe</code>	Probe the Intel Multiprocessor Specification
<code>initrd file</code>	Load an initial ramdisk image file
<code>install options</code>	Install GRUB (deprecated; <code>setup</code> should be used instead)
<code>ioprobe drive</code>	Probe I/O ports used for a drive
<code>kernel file</code>	Load a kernel
<code>lock</code>	Lock a GRUB menu entry
<code>makeactive</code>	Set active partition on root disk to GRUB's root device
<code>map drive1 drive2</code>	Map a drive to another drive
<code>md5crypt</code>	Encrypt a password in MD5 format
<code>module file</code>	Load a kernel module
<code>modulenounzip file</code>	Load a kernel module without decompressing it
<code>pause message</code>	Print a message and wait for a key press
<code>quit</code>	Quit the GRUB shell
<code>reboot</code>	Reboot the system
<code>read address</code>	Read a 32-bit value from memory and print it
<code>root device</code>	Set the current root device
<code>rootnoverify device</code>	Set the current root device without mounting it
<code>savedefault</code>	Save current menu entry as the default entry
<code>setup device</code>	Install GRUB automatically on the device
<code>testload file</code>	Test the filesystem code on a file
<code>testvbe mode</code>	Test a VESA BIOS EXTENSION mode
<code>uppermem kbytes</code>	Set the upper memory size (only for old computers)
<code>vbeprobe mode</code>	Probe a VESA BIOS EXTENSION mode

`dpkg` is the low-level package manager for Debian.
It uses the DEB package format, which is compressed with `ar`.

<code>dpkg -i package.deb</code>	Install a package file
<code>dpkg -r package</code>	Remove a package
<code>dpkg -l</code>	List installed packages and their state
<code>dpkg -L package</code>	List the content of an installed package
<code>dpkg -c package.deb</code>	List the content of a package file
<code>dpkg -S file</code>	Show the package containing a specific file
<code>dpkg-reconfigure package</code>	Reconfigure a package
<code>alien -i package.rpm</code>	Convert an RPM package to DEB and install it. Warning: might break the package database system
<code>dselect</code>	Package manager with text interface, front-end to <code>dpkg</code> . Obsolete

`apt` is the high-level package manager for Debian.

High-level package managers are able to install remote packages and automatically solve dependencies.

<code>apt-get install <i>package</i></code>	Install a package
<code>apt-get remove <i>package</i></code>	Remove a package
<code>apt-get upgrade</code>	Upgrade all installed packages
<code>apt-get dist-upgrade</code>	Upgrade all installed packages and handle dependencies with new versions
<code>apt-get source <i>package</i></code>	Get the source code for a package
<code>apt-get check</code>	Check for broken dependencies and update package cache
<code>apt-get install -f</code>	Fix broken dependencies
<code>apt-get update</code>	Update information on available packages
<code>apt-cache search <i>package</i></code>	Search for a package
<code>apt-cache depends <i>package</i></code>	Show package dependencies
<code>apt-cache show <i>package</i></code>	Show package records
<code>apt-cache showpkg <i>package</i></code>	Show information about a package
<code>apt-file update</code>	Update information about package contents
<code>apt-file list <i>package</i></code>	List the content of an uninstalled package
<code>apt-file search <i>file</i></code>	Show which package provides a specific file
<code>apt-key add <i>keyfile</i></code>	Add a key to the list of keys used to authenticate packages
<code>apt-cdrom add</code>	Add a CD-ROM to the sources list
<code>cat /etc/apt/sources.list</code>	Print list of available repositories
<code>aptitude</code>	Package manager with ncurses UI, front-end to <code>apt</code>
<code>synaptic</code>	Package manager with Gtk+ UI, front-end to <code>apt</code>
<code>katoolin</code>	Install Kali Linux tools and repositories on non-Kali distros. Available for <code>apt</code> -based distros Ubuntu and Linux Mint

`rpm` is the low-level package manager for Red Hat.
It uses the RPM package format, which is `cpio`-compressed.

<code>rpm -i package.rpm</code> <code>rpm -i ftp://host/package.rpm</code> <code>rpm -i http://host/package.rpm</code>	Install a package file
<code>rpm -e package</code>	Remove a package
<code>rpm -U package.rpm</code>	Upgrade a package (and remove old versions)
<code>rpm -F package.rpm</code>	Upgrade a package (only if an old version is already installed)
<code>rpm -qa</code>	List installed packages and their state
<code>rpm -qa --last</code>	List installed packages and their installation date, from newest to oldest
<code>rpm -ql package</code>	List the content of an installed package
<code>rpm -qpl package.rpm</code>	List the content of a package file
<code>rpm -qf file</code>	Show the package containing a specific file
<code>rpm -V package</code>	Verify an installed package
<code>rpm -i package.src.rpm</code>	Install a package source file
<code>rpm -ba package.spec</code>	Compile a package source file
<code>rpm -U --root /path package.rpm</code>	Install a package using an alternative root directory. This is useful in some cases e.g. if the system has been booted from a removable media
<code>rpm2cpio package.rpm</code>	Convert an RPM package to a <code>cpio</code> archive
<code>createrepo directory</code>	Create an XML file of repository metadata from the set of RPMs contained in <i>directory</i>
<code>whohas package</code>	Query multiple package lists to display which version of <i>package</i> is available for different distros
<code>pirut</code>	Package manager with GUI. Obsolete

`yum` (Yellowdog Updater, Modified) was the default high-level package manager from RHEL 5 to RHEL 7. It was based on Yellow Dog Linux's `yup` (Yellowdog Updater). It has been superseded by `dnf` on RHEL 8 and later, but it is still available on these systems, and the two commands retain the same syntax.

<code>yum install <i>package</i></code>	Install a package
<code>yum clean all</code>	Delete temporary files for repositories
<code>rm -rf /var/cache/yum</code>	Delete temporary files for repositories, and free up space taken by orphaned data from disabled, untracked, or removed repos
<code>yumdownloader --resolve <i>package</i></code>	Download package and all its dependencies
<code>yumdownloader --urls <i>package</i></code>	Show URLs that would be downloaded
<code>yum-complete-transaction</code>	Try to complete unfinished or aborted package installations
<code>repoquery --tree-requires <i>package</i></code>	Show a tree with all dependencies of <i>package</i>
<code>needs-restarting</code>	Show a list of process IDs that started running before they, or some component they use, were updated
<code>needs-restarting -r</code>	Report whether a reboot is necessary (exit code 1) or not (exit code 0)
<code>needs-restarting -r echo init 6</code>	Reboot the machine only if necessary

dnf (Dandified YUM) is the default high-level package manager for RHEL 8 and later, and for Fedora. It is a replacement for `yum` while retaining the same syntax.

<code>dnf install package</code>	Install a package
<code>dnf install package.rpm</code>	Install a package file
<code>dnf localinstall package.rpm</code>	
<code>dnf remove package</code>	Remove a package
<code>dnf update package</code>	Update an installed package
<code>dnf update</code>	Update all installed packages
<code>dnf upgrade</code>	Update all installed packages and remove obsolete packages
<code>dnf update --obsoletes</code>	
<code>dnf swap packageout packagein</code>	Replace a package with another
<code>dnf downgrade package</code>	Downgrade an installed package to an older version
<code>dnf list</code>	List all installed and available packages
<code>dnf list searchterm</code>	List installed and available packages matching the search term
<code>dnf list installed</code>	List installed packages
<code>dnf list available</code>	List packages available for install
<code>dnf search searchterm</code>	Search for packages that match the search term in the package name or summary
<code>dnf search all searchterm</code>	Search for packages that match the search term in the package name, summary, or description
<code>dnf deplist package</code>	Show package dependencies (recursively)
<code>dnf list package</code>	Show package records
<code>dnf info package</code>	Show information about a package
<code>dnf history</code>	Show the installation history (installs, updates, etc.)
<code>dnf history list</code>	
<code>dnf history list n</code>	Show item <i>n</i> of the installation history
<code>dnf history info n</code>	Show detailed information on item <i>n</i> of the installation history (begin and end times, packages altered, etc.)
<code>dnf history package package</code>	Show the installation history about a package
<code>dnf history list package package</code>	
<code>dnf whatprovides file</code>	Show which package provides a specific file
<code>dnf provides file</code>	
<code>dnf cmd --disablerepo="*" --enablerepo="repo"</code>	Execute the specified <code>dnf</code> command but only with a specific repository <i>repo</i> enabled
<code>dnf repolist</code>	Print list of available repositories
<code>cat /etc/yum.repos.d/*.repo</code>	
<code>dnf clean all</code>	Delete temporary files for repositories
<code>rm -rf /var/cache/dnf</code>	Delete temporary files for repositories, and free up space taken by orphaned data from disabled, untracked, or removed repos
<code>dnf needs-restarting</code>	Same as the <code>needs-restarting</code> command

Configuration of a yum repository

[fedora]	Repository ID
name=Fedora \$releasever - \$basearch	Repository name
baseurl=http://download.fedoraproject.org/pub/fedora/ linux/releases/\$releasever/Everything/\$basearch/os/ http://foo.org/linux/\$releasever/\$basearch/os/ http://bar.org/linux/\$releasever/\$basearch/os/	List of URLs to the repository's repodata directory. Can be any of these types: file:/// local file file:// NFS http:// HTTP https:// HTTPS ftp:// FTP
enabled=1	Whether this repository is enabled
gpgcheck=1	Whether to perform a GPG signature check on the packages downloaded from this repository
failovermethod=priority	Makes yum try the baseurls in the order they are listed. By default, if more than one baseurl is specified, yum chooses one randomly
metalink=https://mirrors.fedoraproject.org/metalink?repo=fedora-\$releasever&arch=\$basearch	URL to a metalink file that specifies the list of mirrors to use. Can be used with or in alternative to a baseurl
gpgkey=file:///etc/pki/rpm-gpg/ RPM-GPG-KEY-fedora-\$releasever-\$basearch	ASCII-armored GPG public key file of the repository

This repository configuration must be located in a repo file e.g. `/etc/yum.repos.d/fedora.repo`. The same repo file can contain multiple repository definitions.

The manpage `man yum.conf` lists all repository configuration options.

How to install a package on an offline machine

When installing a package on an offline machine, the machine is obviously unable to download the package dependencies. To solve this problem, first create an online machine identical to the offline machine, and with the smallest possible set of packages installed. Then proceed as described below.

On the online machine:

1. Install the package and all its dependencies in a local directory


```
mkdir /tmp/repo
yum --downloadonly --downloadaddir=/tmp/repo install package
```
2. Create a local yum repository from the contents of the local directory


```
createrepo /tmp/repo
chown -R root:root /tmp/repo
chmod -R 755 /tmp/repo
```
3. Transfer the directory `/tmp/repo` from the online machine to the offline machine

On the offline machine:

4. Create a yum repo file `/etc/yum.repos.d/local.repo` for the new repository

```
[local]
name=Local
baseurl=file:///tmp/repo
enabled=1
gpgcheck=0
protect=1
```

5. Install the package from the local repository


```
yum install package
```

`snap` is a distro-independent software packaging and deployment system created by Canonical for systemd systems. Programs are distributed as self-contained packages called **snaps**, and usually made available through the Snap Store. Snaps run in a sandbox, for security reasons. The snap file format is a single compressed SquashFS filesystem.

<code>snap install <i>snap</i></code>	Install a <i>snap</i>
<code>snap info <i>snap</i></code>	Get information about <i>snap</i>
<code>snap find "<i>searchterm</i>"</code>	Find a snap

Flatpak is a distro-independent utility for package management where applications run in a sandboxed environment.

<code>flatpak install <i>repo app</i></code>	Install an application from a remote repository
<code>flatpak install <i>app.flatpak</i></code>	Install an application from a Flatpak file

`opkg` is a lightweight package manager used on embedded Linux systems. It is included in the OpenEmbedded and OpenWrt projects.

<code>opkg install <i>package</i></code>	Install a package
<code>opkg remove <i>package</i></code>	Uninstall a package

`pacman` is Arch Linux's default package manager.

<code>pacman -S <i>package</i></code>	Install a package
<code>pacman -R <i>package</i></code>	Uninstall a package

`npm` is the package manager for the JavaScript runtime environment Node.js.

<code>npm install <i>package</i></code>	Install a package
<code>npm uninstall <i>package</i></code>	Uninstall a package

Homebrew is a package manager developed originally for Apple macOS, but later ported on Linux as well.

<code>brew install <i>package</i></code>	Install a package
<code>brew uninstall <i>package</i></code>	Uninstall a package
<code>brew remove <i>package</i></code>	

`pip` is the Python package installer.

<code>pip install <i>package</i></code>	Install a package
<code>pip uninstall <i>package</i></code>	Uninstall a package

dd	Tool to copy data, byte by byte, from a file or block device. It should not be used on a mounted block device, because of write cache issues
dd if=/dev/sda of=/dev/sdb cat /dev/sda > /dev/sdb	Copy the content of one hard disk over another
dd if=/dev/sda1 of=sda1.img	Generate the image file of a partition
dd if=/dev/cdrom of=cdrom.iso bs=2048	Create an ISO file from a CD-ROM, using a block size transfer of 2 Kb
dd if=install.iso of=/dev/sdc bs=512k	Write an installation ISO file to a device (e.g. a USB thumb drive)
rsync	Tool for local and remote file synchronization. For all copies after the first, copies only the blocks that have changed, making it a very fast and bandwidth-efficient backup solution
rsync -rzv /home /tmp/bak rsync -rzv /home/ /tmp/bak/home	Synchronize the content of the home directory with the temporary backup directory; use recursion, compression, and verbosity
rsync -avz /home root@10.0.0.7:/bak/	Synchronize the content of the home directory with the backup directory on the remote server via SSH; use archive mode (operates recursively and preserves owner, group, permissions, timestamps, and symlinks)
rclone	Tool for cloud storage management. Backups, restores, mirrors, and migrates files from and to a large number of cloud providers (e.g. Ceph, Amazon S3, ownCloud, Google Drive, Microsoft Azure, OpenStack Swift)
burp	Backup and restore program
ddrescue	Data recovery tool. Like dd, but with high tolerance for read errors
testdisk	Data recovery tool. Recovers data from a deleted or corrupted partition
photorec	Data recovery tool. Recovers graphical image and video files from media such as digital cameras and CD-ROMs
ext3grep	Data recovery tool. Recovers deleted files from a EXT3 filesystem
extundelete	Data recovery tool. Recovers deleted files from a EXT3 or EXT4 filesystem
ext4magic	Data recovery tool. Recovers deleted files from a EXT3 or EXT4 filesystem
ntfsundelete	Data recovery tool. Recovers deleted files from a NTFS filesystem
scalpel	Data recovery tool. Recovers data from a disk image or a raw block device

`/dev/st0` First SCSI tape device
`/dev/nst0` First SCSI tape device (no-rewind device file)

`mt` Utility for magnetic tapes
`mt -f /dev/nst0 asf 3` Position the tape at the start of the 3rd file

`mtx` Utility for tape libraries
`mtx -f /dev/sg1 status` Display status of tape library
`mtx -f /dev/sg1 load 3` Load tape from slot 3 to drive 0
`mtx -f /dev/sg1 unload` Unload tape from drive 0 to original slot
`mtx -f /dev/sg1 transfer 3 4` Transfer tape from slot 3 to slot 4
`mtx -f /dev/sg1 inventory` Force robot to rescan all slots and drives
`mtx -f /dev/sg1 inquiry` Inquiry about SCSI media device (Medium Changer = tape library)

cpio	<pre>ls cpio -o > archive.cpio ls cpio -oF archive.cpio find /home/ cpio -o > archive.cpio cpio -id < archive.cpio cpio -i -t < archive.cpio</pre>	Create a cpio archive of all files in the current directory Create a cpio archive of all users' home directories Extract all files, recreating the directory structure List the contents of a cpio archive file
gzip	<pre>gzip file gzip < file > file.gz gunzip file.gz gunzip -tv file.gz zcat file.gz zgrep pattern file.gz zless file.gz zmore file.gz pigz file</pre>	Compress a file with gzip Compress a file with gzip, leaving the original file into place Decompress a gzip-compressed file Test the integrity of a gzip-compressed file Display a gzip-compressed text file grep for a gzip-compressed text file less for a gzip-compressed text file more for a gzip-compressed text file Parallel, multicore-optimized gzip
bzip2	<pre>bzip2 file bunzip2 file.bz2 bzcata file.bz2</pre>	Compress a file with bzip2 Decompress a bzip2-compressed file Display a bzip2-compressed text file
7-Zip	<pre>7z a -t7z archive.7z dir/</pre>	Create a 7-Zip archive
xz	<pre>xz file unxz file.xz xz -d file.xz xzcat file.xz</pre>	Compress a file with xz Decompress a xz-compressed file Display a xz-compressed file
LZMA	<pre>lzma file xz --format=lzma file unlzma file.lzma xz --format=lzma -d file.lzma lzcat file.lzma xz --format=lzma --d --stdout file.lzma</pre>	Compress a file with LZMA Decompress a LZMA-compressed file Display a LZMA-compressed file
rar	<pre>rar a archive.rar dir/ unrar x archive.rar</pre>	Create a RAR archive Extract a RAR archive
tar	<pre>tar cf archive.tar dir/ tar czf archive.tar.gz dir/ tar xzf archive.tar.gz tar cjf archive.tar.bz2 dir/ tar xjf archive.tar.bz2 tar cJf archive.tar.xz dir/ tar xJf archive.tar.xz tar tf archive.tar</pre>	Create a tarred archive Create a tarred gzip-compressed archive Extract a tarred gzip-compressed archive Create a tarred bzip2-compressed archive Extract a tarred bzip2-compressed archive Create a tarred xz-compressed archive Extract a tarred xz-compressed archive List the contents of a tarred archive
star	<pre>star -c -f=archive.star dir/ star -x -f=archive.star</pre>	Create a star archive Extract a star archive
zst	<pre>zstd file unzstd file.zst zstd -d file.zst zstdmt file zstd -T0 file zstdcat file.zst zstd -dcf file.zst</pre>	Compress a file with zst Decompress a zst-compressed file Compress a file with zst, using all physical CPU cores Display a zst-compressed file

<code>man command</code>	Show the manpage (manual page) for <i>command</i>
<code>man n command</code>	Show section <i>n</i> of the <i>command</i> manpage
<code>man man</code>	Show information about manpages' sections: 1 - Executable programs or shell commands 2 - System calls (functions provided by the kernel) 3 - Library calls (functions within program libraries) 4 - Special files 5 - File formats and conventions 6 - Games 7 - Miscellaneous 8 - System administration commands (for root only) 9 - Kernel routines
<code>man n intro</code>	Show an introduction to the contents of section <i>n</i>
<code>mandb</code>	Generate or refresh the search database for manpage entries. This must be done after installing new packages, in order to obtain results from <code>apropos</code> or <code>man -k</code>
<code>whatis command</code>	Show the manpage's short description for <i>command</i>
<code>apropos keyword</code> <code>man -k keyword</code>	Show the commands whose manpage's short description matches <i>keyword</i> . Inverse of the <code>whatis</code> command
<code>apropos -r regex</code> <code>man -k regex</code>	Show the commands whose manpage's short description matches <i>regex</i>
<code>man -K regex</code>	Show the commands whose manpage's full text matches <i>regex</i>
<code>dnf whatprovides /usr/share/man/man/man/ command.n.gz</code> (Red Hat)	Find which package provides section <i>n</i> of the <i>command</i> manpage
<code>dnf install man-pages</code> (Red Hat)	Install a large number of manpages from the Linux Documentation Project
<code>info command</code>	Show the Info documentation for <i>command</i>
<code>catman section</code>	Create or update cat pages , a kind of preformatted manual pages. Obsolete
<code>help</code>	Show the list of available shell commands and functions
<code>help command</code>	Show help about a shell command or function

<code>cat file</code>	Print a text file
<code>cat file1 file2 > file3</code>	Concatenate text files
<code>cat file1 > file2</code> <code>> file2 < file1 cat</code>	Copy <i>file1</i> to <i>file2</i> . The <code>cat</code> command is able to operate on binary streams as well, therefore this command also works on binary files (e.g. JPG images)
<code>cat > file <<EOF</code> <code>line 1</code> <code>line 2</code> <code>line 3</code> <code>EOF</code>	Create a Here Document , storing the lines entered in input to <i>file</i> . <i>EOF</i> can be any text
<code>command <<< 'string'</code>	Create a Here String , passing <i>string</i> as input to <i>command</i>
<code>cat -etv <<< 'string'</code>	Print <i>string</i> , showing all invisible characters
<code>rev <<< 'string'</code>	Reverse <i>string</i>
<code>bat</code>	Improved version of <code>cat</code> with syntax highlighting, auto paging, etc.
<code>tac file</code>	Print or concatenate text files in opposite order line-wise, from last line to first line
<code>rev file</code>	Print a text file with every line reversed character-wise, from last char to first char
<code>column file</code>	Format a text file into columns
<code>pr file</code>	Format a text file for a printer
<code>fmt -w 75 file</code>	Format a text file so that each line has a max width of 75 characters
<code>fold -w40 file</code>	Wrap each line of a text file to 40 characters
<code>nl file</code>	Prepend line numbers to a text file
<code>expand file</code>	Convert tabs into spaces
<code>unexpand file</code>	Convert spaces into tabs
<code>join file1 file2</code>	Join lines of two text files on a common field
<code>paste file1 file2</code>	Merge lines of text files
<code>split -l 1 file</code>	Split a text file into 1-line files named <i>xaa</i> , <i>xab</i> , <i>xac</i> , etc.
<code>sort file</code>	Sort alphabetically the lines of a text file
<code>shuf file</code>	Shuffle randomly the lines of a text file
<code>wc file</code>	Print the number of lines, words, and bytes of a text file
<code>uniq file</code>	Print unique lines of a text file, printing consecutive identical lines only once
<code>uniq -u file</code>	Print lines of a text file that occur only once and are not adjacent to identical lines

<code>head file</code>	Print the first 10 lines of a text file
<code>head -n 10 file</code>	
<code>tail file</code>	Print the last 10 lines of a text file
<code>tail -n 10 file</code>	
<code>tail -f file</code>	Output appended data as the text file grows. Useful to read a logfile in real-time
<code>tail -f file grep --line-buffered pattern</code>	Output appended data as the text file grows, printing only the lines containing <i>pattern</i>
<code>tail -f file stdbuf -o0 grep pattern</code>	
<code>tail -n +1 file1 file2 file3</code>	Print each file with a filename header
<code>multitail -i file1 -i file2</code>	<code>tail</code> for multiple files at the same time (ncurses UI)
<code>cut -d: -f3 file</code>	Cut the lines of a file, considering <code>:</code> as the delimiter and printing only the 3 rd field
<code>cut -d: -f1 /etc/passwd</code>	Print the list of local user accounts in the system
<code>cut -c3-50 file</code>	Print character 3 to 50 of each line of a file
<code>sed 's/foo/bar/' file</code>	Stream Editor: Replace the first occurrence on a line of "foo" with "bar" in <i>file</i> , and print on stdout the result
<code>sed -i 's/foo/bar/' file</code>	Replace "foo" with "bar", overwriting the results in <i>file</i>
<code>sed 's/foo/bar/g' file</code>	Replace all occurrences of "foo" with "bar"
<code>sed '0,/foo/s//bar/' file</code>	Replace only the first line match
<code>sed -n '7,13p' file</code>	Print line 7 to 13 of a text file
<code>sed -n '\$p' file</code>	Print the last line of a text file
<code>sed '\$!d' file</code>	
<code>sed "s/foo/\$var/" file</code>	Replace "foo" with the value of variable <code>\$var</code> . The double quotes are necessary for variable expansion
<code>tr a-z A-Z <file</code>	Translate characters: Convert all lowercase into uppercase in a text file
<code>tr [:lower:] [:upper:] <file</code>	
<code>tr -d 0-9 <file</code>	Delete all digits from a text file
<code>tr -d [:digit:] <file</code>	
<code>tr '[a-zA-Z]' '[n-za-mN-ZA-M]' <file</code>	Perform ROT13 encoding (i.e. decoding) of a text file
<code>rp1 oldstring newstring file</code>	Replace strings in a file
<code>awk</code>	Interpreter for the AWK programming language, designed for text processing and data extraction
<code>gawk</code>	GNU's version of AWK interpreter
<code>pgawk</code>	GNU's version of AWK interpreter, with profiling

<code>grep foo file</code>	Print the lines of a file containing "foo"
<code>grep -v foo file</code> <code>grep --invert-match foo file</code>	Print the lines of a file not containing "foo"
<code>grep -e foo -e bar file</code> <code>grep -E 'foo bar' file</code>	Print the lines of a file containing "foo" or "bar"
<code>grep -v -e foo -e bar file</code>	Print the lines of a file containing neither "foo" nor "bar"
<code>grep -l foo *.*</code>	List the files that have at least one line containing "foo"
<code>grep -L foo *.*</code> <code>grep --files-without-matches foo *.*</code>	List the files where none of the lines contain "foo"
<code>grep -E regex file</code> <code>egrep regex file</code>	Print the lines of a file matching the given Extended Regex
 <code>look string file</code>	 Display lines from <i>file</i> (which must be sorted) beginning with <i>string</i> . If <i>file</i> is not specified, <code>/usr/share/dict/words</code> is used
 <code>diff file1 file2</code> <code>diff -u file1 file2 diff-so-fancy</code>	 Compare two text files line by line and print the differences Improved version of <code>diff</code> with better readability
 <code>comm file1 file2</code>	 Compare two sorted files line by line, and print three columns: lines unique to <i>file1</i> , lines unique to <i>file2</i> , and lines common to both files
 <code>strings file</code>	 Show all printable character sequences at least 4-characters long contained in <i>file</i>
 <code>comm -12 <(sort file1) <(sort file2)</code> <code>grep -F -x -f file2 file1</code>	 Print the lines from <i>file1</i> that are also present in <i>file2</i>

<code>uuencode file</code>	Encode a file using only printing ASCII characters. Used to transmit a binary file over a medium that supports only ASCII data, e.g. e-mail
<code>uuencode -m file</code>	Encode a file to Base64
<code>uudecode file</code>	Decode a file
<code>base64 file</code>	Encode a file to Base64
<code>base64 -d file</code>	Decode a file from Base64
<code>od file</code>	Dump a file into octal (or other formats)
<code>hexdump options file</code> <code>hd options file</code>	Dump a file into hexadecimal (or octal, decimal, ASCII)
<code>hexedit</code> <code>hexeditor</code>	Hexadecimal editor
<code>xxd options file</code>	Convert a file from binary to hexadecimal (i.e create a hex dump), or vice versa
<code>cmp file1 file2</code>	Compare two files byte by byte and print the differences. Similar to <code>diff</code> , but it operates on binary files instead
<code>tidy</code>	Correct and tidy up the markup of HTML, XHTML, and XML files
<code>tidy -asxml -xml -indent -wrap 2000 \</code> <code>-quiet --hide-comments yes file.xml</code>	Strip out comments from an XML file
<code>json_verify < file.json</code>	Validate the syntax of a JSON file
<code>json_reformat < file.json</code>	Pretty format a JSON file
<code>jq</code>	JSON processor
<code>fx</code>	JSON viewer and processing tool
<code>pandoc options file</code>	Convert a file from a markup format to another (e.g. HTML, LaTeX, RTF, Markdown, OpenDocument XML, ODT, DOCX, PDF, EPUB)
<code>antiword file.doc</code>	Show text and images from a MS Word document
<code>catdoc file.doc</code>	Output plaintext from a MS Word document

<code>^</code>	Beginning of a line
<code>\$</code>	End of a line
<code>\< \></code>	Word boundaries (beginning of line, end of line, space, or punctuation mark)
<code>.</code>	Any character except newline
<code>[abc]</code>	Any of the characters specified
<code>[a-z]</code>	Any of the characters in the specified range
<code>[^abc]</code>	Any character except those specified
<code>*</code>	Zero or more times the preceding regex
<code>+</code>	One or more times the preceding regex
<code>?</code>	Zero or one time the preceding regex
<code>{5}</code>	Exactly 5 times the preceding regex
<code>{5,}</code>	5 times or more the preceding regex
<code>{,10}</code>	At most 10 times the preceding regex
<code>{5,10}</code>	Between 5 and 10 times the preceding regex
<code> </code>	The regex either before or after the vertical bar
<code>()</code>	Grouping, to be used for back-references. <code>\1</code> expands to the 1 st match, <code>\2</code> to the 2 nd , etc. up to <code>\9</code>

The symbols above are used in POSIX EREs (Extended Regular Expressions).

In POSIX BREs (Basic Regular Expressions), the symbols `?` `+` `{` `|` `( )` need to be escaped; this is done by adding a backslash character `\` in front of them.

<code>cp file file2</code>	Copy a file
<code>cp file dir/</code>	Copy a file to a directory
<code>cp -ar /dir1/. /dir2/</code>	Copy a directory recursively
<code>mv file file2</code>	Rename a file
<code>mv file dir/</code>	Move a file to a directory
<code>rm file</code>	Delete a file
<code>rename str1 str2 file</code>	Rename a file, replacing in the filename the first occurrence of string <i>str1</i> with <i>str2</i>
<code>rename .htm .html *.htm</code>	Rename all <i>.htm</i> files to <i>.html</i>
<code>mmv src dest</code>	Mass-move multiple files by wildcard pattern
<code>mmv '*.htm' '#1.html'</code>	Rename all <i>.htm</i> files to <i>.html</i>
<code>mcp src dest</code>	Mass-copy multiple files by wildcard pattern
<code>mad src dest</code>	Mass-append content from multiple files by wildcard pattern
<code>mln src dest</code>	Mass-link multiple files by wildcard pattern
<code>unlink file</code>	Remove the hard link to a file. Equivalent to <code>rm</code>
<code>touch file</code>	Change access timestamp and modify timestamp of a file as now. If the file does not exist, it is created
<code>truncate -s size file</code>	Shrink or extend a file to the specified size. If the file is larger than the specified size, it is truncated; if the file is shorter, the extra space is filled with zeros
<code>fallocate -l 1G file</code>	Preallocate 1 Gb of space for a file
<code>mktemp</code>	Create a temporary file or directory, using <code>tmp.XXXXXXXXXX</code> as filename template
<code>fdupes dir</code>	Examine a directory for duplicate files in it. To consider files a duplicate, it first compares file sizes and MD5 signatures, then file contents byte-by-byte
<code>shred /dev/hda</code>	Securely wipe the contents of a device
<code>shred -u file</code>	Securely delete a file

File-naming wildcards (globbing)

<code>*</code>	Matches zero or more characters
<code>?</code>	Matches one character
<code>[abc]</code>	Matches a, b, or c
<code>[!abc]</code>	Matches any character except a, b, or c
<code>[a-z]</code>	Matches any character between a and z

Brace expansion

<code>cp foo.{txt,bak}</code>	Copy file "foo.txt" to "foo.bak"
<code>touch foo_{a,b,c}</code>	Create files "foo_a", "foo_b", "foo_c"
<code>touch foo_{a..c}</code>	

<code>cd <i>directory</i></code>	Change to the specified directory
<code>cd -</code>	Change to the previously used directory
<code>pwd</code>	Print the current working directory
<code>ls</code> <code>dir</code> <code>vdir</code>	List the contents of the current directory
<code>ls -d */</code>	List only directories contained on the current directory
<code>ls -lap --sort=v</code>	List files, sorted by version number
<code>mkdir <i>dir</i></code>	Create a directory
<code>mkdir -m 755 <i>dir</i></code>	Create a directory with mode 755
<code>mkdir -p /<i>dir1</i>/<i>dir2</i>/<i>dir3</i></code>	Create a directory, creating also the parent directories if they don't exist
<code>rmdir <i>dir</i></code>	Delete a directory (which must be empty)
<code>tree</code>	List directories and their contents in hierarchical format
<code>dirs</code>	Display the directory stack (i.e. the list of remembered directories)
<code>pushd <i>dir</i></code>	Add <i>dir</i> to the top of the directory stack and make it the current working directory
<code>popd</code>	Remove the top directory from the directory stack and change to the new top directory
<code>dirname <i>file</i></code>	Output the directory path in which <i>file</i> is located, stripping any non-directory suffix from the filename
<code>realpath <i>file</i></code>	Output the resolved absolute path of <i>file</i>

Bash directory shortcuts	
<code>.</code>	Current directory
<code>..</code>	Parent directory
<code>~</code>	Home directory of current user
<code>~<i>user</i></code>	Home directory of <i>user</i>
<code>~-</code>	Previously used directory

<code>lsuf</code>	List all open files
<code>lsuf -u user</code>	List all files currently open by <i>user</i>
<code>lsuf -i</code>	List open files and their sockets (equivalent to <code>netstat -ap</code>)
<code>lsuf -i :80</code>	List connections of local processes on port 80
<code>lsuf -iTCP:70-90</code>	List connections of local processes on TCP ports between 70 and 90
<code>lsuf -i@10.0.0.3</code>	List connections of local processes to remote host 10.0.0.3
<code>lsuf -i@10.0.0.3:80</code>	List connections of local processes to remote host 10.0.0.3 on port 80
<code>lsuf -c mysqld</code>	List all files opened by <code>mysqld</code> , the MySQL daemon
<code>lsuf file</code>	List all processes using a specific <i>file</i>
<code>lsuf +L1</code>	List open files with a link count smaller than 1 i.e. that have been unlinked. These files are not accessible anymore but take up disk space. A process holding such a file prevents the system from deleting it (thus freeing disk space), until the process is killed or restarted
<code>fuser</code>	Show the name of processes using a specific file, directory, or socket
<code>fuser -v file</code>	Show the name of the process using <i>file</i>
<code>fuser -v -n tcp 443</code>	Show the name of the process running on port 443
<code>lslocks</code>	List information about all currently held file locks
<code>lslk</code>	List information about all locks currently held on files with local inodes
<code>flock</code>	Manage file locks
<code>tmpwatch</code>	Remove files which have not been accessed for some time
<code>stat file</code>	Display file or filesystem status
<code>stat -c %A file</code>	Display file permissions
<code>stat -c %s file</code>	Display file size, in bytes
<code>crc32 file</code>	Calculate the CRC-32 checksum of <i>file</i> . As the CRC-32 checksum is not a cryptographic hash, it must only be used for error detection in transmission and storage, and not to detect malicious modifications to the file
<code>namei file</code>	Follow the path to <i>file</i> (a symlink, file, directory, etc.) until a terminal point is found, and then display its type (file, directory, char device, etc.)
<code>file file</code>	Analyze the content of a file or directory, and display the kind of file (e.g. executable, text file, program text, swap file)

In Linux, everything is (displayed as) a file. **File descriptors** are unique identifiers for any I/O resource e.g. a file, pipe, or network socket; they are automatically associated to any process launched.

Standard POSIX file descriptors				
File descriptor #	Name	Type	Default device	Device file
0	Standard input (stdin)	Input text stream	Keyboard	/dev/stdin
1	Standard output (stdout)	Output text stream	Terminal	/dev/stdout
2	Standard error (stderr)	Output text stream	Terminal	/dev/stderr

<code>mail user@email < file</code>	Redirect <i>file</i> to the stdin of command <code>mail</code> (in this case, send via e-mail the contents of <i>file</i> to the email address <i>user@email</i>). Redirection is handled by the shell, not by the command invoked. The space after the redirection operator can be omitted
<code>ls > file</code> <code>ls 1> file</code>	Redirect the stdout of command <code>ls</code> to <i>file</i> (in this case, write on <i>file</i> the contents of the current directory). This overwrites <i>file</i> if it already exists, unless the Bash noclobber option is set (via <code>set -o noclobber</code>)
<code>ls > file</code>	Redirect the stdout of command <code>ls</code> to <i>file</i> , even if noclobber is set
<code>ls >> file</code> <code>ls 1>> file</code>	Append the stdout of command <code>ls</code> to <i>file</i>
<code>ls 2> file</code>	Redirect the stderr of command <code>ls</code> to <i>file</i> (in this case, write any error encountered by the command <code>ls</code> to <i>file</i>)
<code>ls 2>> file</code>	Append the stderr of command <code>ls</code> to <i>file</i>
<code>ls 2> /dev/null</code>	Silence any error coming from the command <code>ls</code>
<code>cat <file1 >file2</code> <code><file1 cat >file2</code> <code><file1 >file2 cat</code>	Redirect <i>file1</i> to the stdin and <i>file2</i> to the stdout of the command <code>cat</code> (in this case, copy <i>file1</i> to <i>file2</i>). <code>cat >file2 <file1</code> also works, but is not recommended because it truncates <i>file2</i> if for any reason <i>file1</i> cannot be opened
<code>cat /etc/passwd wc -l</code>	Pipe the stdout of command <code>cat</code> to the stdin of command <code>wc</code> (in this case, print the number of accounts in the system). Piped commands run concurrently
<code>echo "\$(sort file)" >file</code> <code>echo "`sort file`" >file</code> <code>sort file sponge file</code>	Sort the contents of <i>file</i> and write the output to the file itself. <code>sort file > file</code> would not produce the desired result, because the shell would create the stdout destination (deleting the content of the preexisting <i>file</i>) before running the <code>sort</code> command
<code>ls 2>&1</code> <code>ls >file 2>&1</code>	Redirect stderr of command <code>ls</code> to stdout Redirect both stdout and stderr of command <code>ls</code> to <i>file</i> . Commands <code>ls &> file</code> and <code>ls >& file</code> also work on some systems, but are not recommended because they are not POSIX standard
<code>>file</code>	Create an empty file. If the file exists, its content will be deleted
<code>mkfifo namedpipe</code>	Create a FIFO aka named pipe. The commands <code>ls ~ >namedpipe</code> and <code>cat namedpipe</code> (run in different terminals, as the writing on the named pipe is blocking until the content of the FIFO is consumed) will then display the contents of the home directory

<code>tee file</code>	Read from stdin and write both to stdout and <i>file</i>
<code>tee -a file</code>	Read from stdin and append both to stdout and <i>file</i>
<code>ls tee file</code>	Write the contents of the current directory to screen and to <i>file</i> at the same time
<code>stdbuf option command</code>	Run <i>command</i> with modified stdin, stdout, or stderr buffering
<code>sponge file</code>	Read from stdin and write to <i>file</i> , absorbing all input before opening the output <i>file</i> for writing
<code>sort file sponge file</code>	Sort the contents of <i>file</i> and write the output to the file itself
<code>ifne</code>	Run a command only if stdout is not empty
<code>find . -name file ifne mail root</code>	If there is <i>file</i> in the current directory, send it via e-mail to the root user
<code>pv</code>	Monitor the progress of data through a pipe
<code>pv file > file2</code>	Copy a file, showing a progress meter
<code>pv -qL10 <<< "message"</code>	Print <i>message</i> onscreen, one character at a time

<pre>read MYVAR</pre>	Read a variable from standard input
<pre>read -n 8 MYVAR</pre>	Read only max 8 characters from standard input
<pre>read -t 60 MYVAR</pre>	Read a variable from standard input, timing out after one minute
<pre>read -s MYVAR</pre>	Read a variable from standard input without echoing to terminal (silent mode)
<pre>while read -r line do echo "Hello \$line" done < file</pre>	Process a text file line by line, reading from <i>file</i> , and output the lines. If <i>file</i> is <code>/dev/stdin</code> , reads from standard input instead
<pre>while read line do for word in \$line do echo "Hello \$word" done done < file</pre>	Process a text file containing multiple words in each line, and output the words
<pre>while IFS=\$'\t' read -r -a array do echo "\${array[0]}" echo "\${array[1]}" echo "\${array[2]}" done < file</pre>	Process a text file containing three words per line separated by a tab, and output the words. Example of input file: <pre>aaaa bbb ccc dd eeeee ff ggg hhh iiii</pre>
<pre>readarray MYARRAY mapfile MYARRAY</pre>	Read lines from stdin into an indexed array
<pre>echo \$MYVAR</pre>	Print a variable on screen
<pre>echo -n "message" printf "message"</pre>	Print <i>message</i> onscreen without a trailing line feed
<pre>echo -e '\a'</pre>	Produce an alert sound (BEL sequence)
<pre>echo .*</pre>	Resolve globs, printing all files whose name begins with a dot in the current dir
<pre>echo rm -f .*</pre>	Resolve globs, expanding the filenames and printing the actual <code>rm</code> command that would have been executed

Any application, program, script, or service that runs on the system is a **process**. Segments of a process are called **threads**, which are lightweight and share the same memory space.

The kernel uses **signals** to communicate with a process. Signals are also used for inter-process communication (however, it is always the kernel which verifies whether a process is authorized to send a specific signal to another process).

Each process has a unique **PID (Process ID)** and a **PPID (Parent Process ID)**: when a process spawns a child, the process PID is assigned to the child's PPID.

The process with PID 1 (`init` or `systemd`) is the ancestor of all processes and is unkillable; its death causes a kernel panic. When a process terminates, the system assigns a new parent to its orphaned child. This new parent process is the nearest ancestor process of the child that marked itself as a subreaper, or the process with PID 1 if there is no such ancestor subreaper process.

A **zombie** process is a process that has terminated execution but whose parent, for some reason, failed to reap. When a child process dies, its status becomes `EXIT_ZOMBIE` and a `SIGCHLD` is sent to the parent. The parent should then call the `wait()` system call to read the dead process' exit status and other information; until that moment, the child process remains a zombie.

Zombie processes do not take up system resources and are usually not a problem, but may be a symptom of sloppy programming in the parent program's code or of system bugs. To eliminate a zombie, its parent must be terminated by sending it a `SIGKILL`; the orphaned zombie will then be reaped by the nearest subreaper process or by process with PID 1. It is also possible to send a `SIGCHLD` to the parent to ask it to reap the child, but this is not guaranteed to work because the parent might ignore (or not have a proper signal handler for) `SIGCHLD` signals altogether.

cgroups (control groups) are a feature of the Linux kernel allowing the organization of processes into hierarchical groups for monitoring and rate limiting purposes. Many projects (e.g. `systemd`, `Docker`, and `Kubernetes`) use it.

```
ps -ef (UNIX options)    List all processes
ps aux (BSD options)
```

```
pstree PID              Display all processes in hierarchical format.
                        The process tree is rooted at PID, or at init if PID is omitted
```

```
pidof processname       Show PIDs of processes with name processname
pidof -s processname     Show PID of process with name processname, returning a single result
```

```
pgrep sshd              Show processes whose name is "sshd"
ps -ef | grep "[s]shd"
pgrep -u root sshd       Show processes whose name is "sshd" and are owned by root
```

```
pmap PID                Display the memory map of process PID
```

```
:(){ :|:& };:            Fork bomb: starts a process that continually replicates itself, slowing down or crashing the
                        system because of resource starvation. Dangerous!
```

```
( command )& pid=$!; \   Run command and kill it after n seconds
sleep n; kill -9 $pid
```

To each process is associated a **niceness** value: the higher the niceness, the lower the priority. The niceness value ranges from -20 to 19, and a newly created process has a default niceness of 0. Unprivileged users can modify a process' niceness only within the range from 1 to 19.

```
nice -n -5 command      Start command with a niceness of -5. If niceness is omitted, a default value of 10 is used
renice -5 command        Change the niceness of a running command to -5
snice                    Change the niceness of a process. Obsolete
```

Processes whose parent is a shell are called **jobs**.

<code>jobs</code>	List all jobs
<code>CTRL Z</code>	Suspend a job, putting it in the stopped state (send a SIGTSTP)
<code>bg %n</code>	Put job <i>#n</i> in the background (send a SIGCONT)
<code>fg %n</code>	Resume job <i>#n</i> in the foreground and make it the current job (send a SIGCONT)
<code>kill %n</code>	Kill job <i>#n</i> (send a SIGKILL)
<code>disown %n</code>	Remove job <i>#n</i> from the table of active jobs
<code>disown -h %n</code>	Prevent job <i>#n</i> from receiving a SIGHUP if the shell receives that signal
<code>nohup script.sh</code>	Prevent a process from terminating (receiving a SIGHUP) when its parent Bash dies. When a Bash shell is terminated cleanly via <code>exit</code> , its jobs become child of the Bash's parent and continue running. When a Bash shell is killed instead, it issues a SIGHUP to its children which terminate execution

Most common signals		
Signal number	Signal name	Meaning
1	SIGHUP	Reload configuration
2	SIGINT	Interrupt execution
3	SIGQUIT	Interrupt execution and produce a core dump. User-initiated
4	SIGILL	Illegal instruction
5	SIGTRAP	Debugger breakpoint
6	SIGABRT	Abort the program due to a critical error
7	SIGBUS	Bus error. Also caused by any general device fault
8	SIGFPE	Floating point exception, or any other fatal arithmetic error
9	SIGKILL	Kill unconditionally. This signal cannot be ignored
10	SIGUSR1	User-defined signal 1
11	SIGSEGV	Segmentation fault
12	SIGUSR2	User-defined signal 2
13	SIGPIPE	Attempt to write to a pipe without a process connected to the other end
14	SIGALRM	A time interval specified in a call to the alarm function has expired
15	SIGTERM	Terminate gracefully
16	SIGSTKFLT	Stack fault on coprocessor
17	SIGCHLD	Child is stopped or terminated, request to reap it
18	SIGCONT	Continue execution
19	SIGSTOP	Stop execution. This signal cannot be ignored
20	SIGTSTP	Stop execution

The manpage `man 7 signal` lists all possible signal numbers and names.

<code>kill -l</code>	List all available signal names
<code>kill -l <i>n</i></code>	Print the name of signal number <i>n</i>
<code>kill -<i>signumber</i> <i>PID</i></code> <code>kill -s <i>SIGNAME</i> <i>PID</i></code>	Send a signal to process <i>PID</i>
<code>kill -9 1138</code>	Send a signal 9 (SIGKILL) to process 1138, hence killing it
<code>kill -s SIGCHLD <i>PPID</i></code>	Eliminate a zombie process by sending its parent (<i>PPID</i>) a SIGCHLD
<code>killall -9 sshd</code>	Kill processes whose name is "sshd"
<code>pkill -9 -u root sshd</code>	Kill processes whose name is "sshd" and are owned by root
<code>pkill -9 -u user</code>	Kill all processes owned by <i>user</i> , forcing the user to log out
<code>skill</code>	Send a signal to a process or show process status. Obsolete
<code>xkill</code>	Kill a process by its X GUI resource. Pops up a cursor to select a window
<code>trap <i>action condition</i></code> <code>trap myfunc EXIT</code>	Trap a signal In a script, execute the function <code>myfunc</code> whenever the script exits (whether because it reaches its end naturally, or it encounters an <code>exit</code> command, or it receives a SIGINT termination signal, or a fatal error occurs, etc.)
<code>strace <i>command</i></code>	Trace the execution of <i>command</i> , intercepting and printing system calls called by a process and signals received by a process
<code>ipcs</code>	Show IPC (inter-process communication) facilities information: shared memory, message queues, and semaphores
<code>forkstat</code>	Analyze process activity with respect to <code>fork()</code> , <code>exec()</code> , and <code>exit()</code> system calls. Useful to monitor system behaviour and track down rogue processes

There are many full-fledged resource monitoring solutions for Linux environments: some of them are Munin, Zabbix, Centreon, and Nagios (system and network monitor and alert tools), MRTG and Cacti (network monitors), and Netdata (real-time performance and health monitor).

<code>top</code>	Monitor processes in real-time
<code>htop</code>	Monitor processes in real-time (ncurses UI)
<code>iostat</code>	Display I/O usage by processes in the system
<code>atop</code>	Advanced system monitor that displays the load on CPU, RAM, disk, and network
<code>powertop</code>	Power consumption and power management diagnosis tool
<code>nvtop</code>	Neat Videocard top, a task monitor for GPUs (AMD, NVIDIA, etc.)
<code>uptime</code>	Show how long the system has been up, how many users are connected, and the system load averages for the past 1, 5, and 15 minutes
<code>sysbench</code>	Multi-threaded benchmark tool able to monitor different OS parameters: file I/O, scheduler, memory allocation, thread implementation, databases
<code>inxi</code>	Debugging tool to rapidly and easily gather system information and configuration
<code>stress-ng</code>	Tool for CPU and RAM stress tests
<code>collectd</code>	System statistics collector
<code>turbostat</code>	Tool to report processor topology, frequency, idle power-state statistics, temperature, and power
<code>sensors</code>	Print sensor chips information (e.g. temperature)
<code>psensor</code>	GUI client tool for monitoring hardware sensors (e.g. temperature, fan speed) of a remote server
<code>psensor-server</code>	HTTP server for <code>psensor</code>
<code>corefreqd</code>	Daemon for CoreFreq, a CPU monitoring tool with BIOS-like functionalities
<code>corefreq-cli</code>	CoreFreq client
<code>sysmon</code>	Monitor for system events. Developed by Sysinternals
<code>conky</code>	System monitor widget GUI with integration for audio player, email, and news
<code>gkrellm</code>	System monitor widget GUI

vmstat	Print a report about virtual memory statistics: processes, memory, paging, block I/O, traps, disks, and CPU activity
iostat	Print a report about CPU utilization, device utilization, and network filesystem. The first report shows statistics since the system boot; subsequent reports will show statistics since the previous report
mpstat	Print a report about processor activities
vmstat <i>n m</i>	Print the relevant report every <i>n</i> seconds for <i>m</i> times
iostat <i>n m</i>	
mpstat <i>n m</i>	

Output of command `vmstat`

```
procs -----memory----- --swap-- ----io---- --system-- -----cpu-----
r  b  swpd  free  buff  cache  si  so  bi  bo  in  cs  us  sy  id  wa  st
0  0      0 296724 267120 3393400  0  0  17  56  0  3  2  2 95  1  0
```

procs	r	Number of runnable processes (running or waiting for run time)
	b	Number of processes in uninterruptible sleep
memory	swpd	Virtual memory used (swap)
	free	Free memory (idle)
	buff	Memory used as buffers in Kb
	cache	Memory used as cache
swap	si	Memory swapped in from disk
	so	Memory swapped out to disk in Kb/second
io	bi	Blocks received in from a block device
	bo	Blocks sent out to a block device in blocks/second
system	in	Number of interrupts per second
	cs	Number of context switches
cpu	us	Time spent running user code (non-kernel)
	sy	Time spent running system code (kernel)
	id	Time spent idle in percentage of total CPU time
	wa	Time spent waiting for I/O
	st	Time stolen from a virtual machine

free

Show the amount of free and used memory in the system

Output of command <code>free</code>						
	total	used	free	shared	buff/cache	available
Mem:	16344088	2273312	11531400	776228	2539376	12935112
Swap:	1048572	0	1048572			
	total	used	free	shared	buffers	cached
Mem:	1504544	1491098	13021	0	91112	764542
-/+ buffers/cache:		635212	869498			
Swap:	2047686	7667	2040019			

Mem	total	Total configured amount of memory
	used	Used memory
	free	Unused memory
	shared	Memory used by tmpfs, 0 if not available
	buff/cache	Memory used by kernel buffers, page cache, and slabs
	available	Memory available for new applications (without using swap) *
-/+ buffers/cache	used	Memory used by kernel buffers
	free	Memory available for new applications (without using swap) *
Swap	total	Total configured amount of swap space
	used	Used swap space
	free	Free swap space *

* These are the true values indicating the free system resources available.

All values are in kB, unless unit options are used:

- b Show values in bytes
- k Show values in kilobytes (kB)
- m Show values in megabytes (MB)
- g Show values in gigabytes (GB)
- h Show values in human-readable format, i.e. scaled to the shortest 3-digit unit and displaying the unit
- si Show values according to the International Systems of Units (powers of 1000 instead of powers of 1024)

PCP (Performance Co-Pilot) is an open source framework and toolkit for monitoring and analyzing system performance, either live or historical.


<code>pcp</code>	Run a PCP command or summarize an installation
<code>pminfo</code>	Display information about performance metrics
<code>pmdumptext</code>	Dump performance metrics to a text table
<code>pmrep</code>	Performance metrics reporter
<code>pmstat</code>	Periodically print a one-line summary of system performance at the highest level
<code>pmval</code> <code>pmevent</code>	Print values of a performance metric

rasdaemon is a RAS (Reliability, Availability, and Serviceability) logging tool that records memory and disk errors. It uses the Linux kernel's EDAC (Error Detection and Correction) subsystem which gathers data from those computer components which have error-checking capabilities enabled.

<code>rasdaemon --enable</code>	Enable RAS tracing events
<code>rasdaemon --record</code>	Record RAS events permanently via a SQLite3 database
<code>ras-mc-ctl</code>	RAS administration tool
<code>ras-mc-ctl --summary</code>	Print a summary of the recorded errors
<code>ras-mc-ctl --errors</code>	Show the recorded errors

sar (System Activity Report) is a tool to collect, report, or save system activity information. Reports are generated from data collected via the cron job `sysstat` and stored in `/var/log/sa/sa $n$` , where n is the day of the month.


<code>sar</code>	Show reports about system activity, including reboots
<code>sar -f /var/log/sa/sa7 -s 06:00:00 -e 09:00:00</code>	Show reports for activity from 6 to 9 AM on the 7 th of the month
<code>sar -u n m</code>	Show real-time CPU activity, every n seconds for m times
<code>sar -n DEV</code>	Show real-time network activity (received and transmitted packets per second)



Permission	Octal value	Command	Effect on file	Effect on directory
Read	user: 400	chmod u+r	Can open and read the file	Can list directory content
	group: 40	chmod g+r		
	others: 4	chmod o+r		
Write	user: 200	chmod u+w	Can modify the file	Can create, delete, and rename files in the directory
	group: 20	chmod g+w		
	others: 2	chmod o+w		
Execute	user: 100	chmod u+x	Can execute the file (binary or script)	Can enter the directory, and search files within (by accessing a file's inode)
	group: 10	chmod g+x		
	others: 1	chmod o+x		
SetUID (SUID)	4000	chmod u+s	Executable is run with the privileges of the file's owner	No effect
SetGID (SGID)	2000	chmod g+s	Executable is run with the privileges of the file's group	All new files and subdirectories inherit the directory's group ID
Sticky	1000	chmod +t	No effect. (Historically: keep file in main memory)	Files inside the directory can be deleted or moved only by the file's owner

File **permissions** (aka file **mode**) are not cumulative; the most specific permission (i.e. user vs group, and group vs others) applies.

<code>chmod 711 file</code> <code>chmod u=rwx,go=x file</code>	Set read, write, and execute permission to user; set execute permission to group and others
<code>chmod u+wx file</code>	Add write and execute permission to user
<code>chmod -x file</code>	Remove execute permission from everybody (user, group, and others)
<code>chmod -R g+x /path</code>	Set the group execute bit recursively on <i>path</i> and every directory and file underneath
<code>find /path -type d \</code> <code>-exec chmod g+x {} \;</code>	Set the group execute bit recursively on <i>path</i> and every directory, but not file, underneath
<code>chown user file</code>	Change the owner of the file to <i>user</i>
<code>chown user:group file</code>	Change the owner of the file to <i>user</i> , and group ownership of the file to <i>group</i>
<code>chown :group file</code> <code>chgrp group file</code>	Change group ownership of the file to <i>group</i>
<code>umask 022</code>	Set the permission mask to 022, hence masking write permission for group and others. Linux default permissions are 0666 for files and 0777 for directories. These base permissions are ANDed with the inverted umask value to calculate the final permissions of a newly created file or directory



regular file	A regular file. Passes the <code>test -f</code> test.
directory	A directory. Passes the <code>test -d</code> test.
symbolic link	A symlink. Created via the <code>ln -s</code> command.
named pipe aka FIFO	An extension of the traditional unnamed UNIX pipe, which terminates when the interconnecting processes end. Instead, a named pipe can last beyond the life of the processes, until it is deleted. Facilitates IPC (inter-process communication) by allowing local programs to exchange data. As in unnamed pipes, communication is unidirectional, from the stdout of the first process to the stdin of the second. Created via the <code>mkfifo</code> command.
UNIX domain socket	A pseudo file that represents a network connection. A socket can be bound either to an "internet address and port" tuple (AF_INET) or to a special file (AF_UNIX). Once a network socket is created, writes to the socket are turned into network packets that get sent out, and data received from the network can be read from the socket. Facilitates inter-process communication by allowing a local program to exchange data bidirectionally with a remote program.
character device file	Interface to a device driver. Provides unbuffered, direct access to the device as a serial stream of input/output.
block device file	Interface to a device driver. Provides buffered, random access. Allows reading/writing a block of any size.

<code>chattr +attribute file</code>	Add a file or directory attribute
<code>chattr -attribute file</code>	Remove a file or directory attribute
<code>chattr =attribute file</code>	Set a file or directory attribute, removing all other attributes
 <code>lsattr file</code>	 List file or directory attributes

Attribute	Effect
a	The file can only be opened in append mode for writing
A	When the file is accessed, its atime record is not modified; this reduces I/O operations, improving performances. Equivalent to the <code>noatime</code> mount option
c	The file is automatically compressed on the fly on disk by the kernel; reading it produces uncompressed data
C	For filesystems which perform copy-on-write, the file is not subject to copy-on-write updates
d	The file will not be backed up by the <code>dump</code> program
D	When the directory is modified, changes are written synchronously on disk. Equivalent to the <code>dirsync</code> mount option
e	The file is using extents for mapping the blocks on disk
E	Compression error on a compressed file. This attribute is used by experimental compression patches
h	The file stores its blocks in units of filesystem blocksize instead of in units of sectors. This means that the file is (or was) larger than 2 Tb
i	The file is immutable i.e. cannot be deleted, modified, renamed, linked, or changed permissions
I	The directory is being indexed using hashed trees
j	All file data is written to the ext3 or ext4 journal before being written to the file itself
N	The file has data stored inline within the inode itself
s	The file will be securely wiped by zeroing when deleted
S	When the file is modified, changes are written synchronously on disk. Equivalent to the <code>sync</code> mount option
t	For filesystems with support for tail-merging, the file will not have EOF partial block fragment merged with other files. This is necessary for some applications e.g. LILO
T	The directory is the top of directory hierarchies for the purpose of the Orlov block allocator
u	The file can be undeleted after being deleted
X	Raw contents of the compressed file are directly accessible. This attribute is used by experimental compression patches
Z	The compressed file is dirty. This attribute is used by experimental compression patches

Timestamp	Value tracked	Displayed via
mtime	Time of last modification to file contents (data itself)	<code>ls -l</code>
ctime	Time of last change to file contents or file metadata (owner, group, permissions, link count, etc. not including timestamps)	<code>ls -lc</code>
atime	Time of last access to file for reading contents	<code>ls -lu</code>
btime	Time of file birth/creation . Not a POSIX standard; available only for those filesystems (e.g. ext4, XFS, JFS, Btrfs) that store this value	<code>ls -lt</code>

Access Control Lists (ACLs) provide a fine-grained set of permissions that can be applied to files and directories. An **access ACL** is set on an individual file or directory; a **default ACL** is set on a directory, and applies to all files and subdirectories created inside it that don't have an access ACL. The final permissions are the intersection of the ACL with the `chmod/umask` value. A partition must have been mounted with the `acl` option in order to support ACLs on files.

<code>setfacl -m u:user:permissions file</code>	Set an access ACL on a file for a user
<code>setfacl -m g:group:permissions file</code>	Set an access ACL on a file for a group
<code>setfacl -m m:permissions file</code>	Set the effective rights mask on a file
<code>setfacl -m o:permissions file</code>	Set the permissions on a file for other users
<code>setfacl -x u:user file</code>	Remove an access ACL from a file for a user
<code>setfacl -x g:group file</code>	Remove an access ACL from a file for a group

The *permissions* are standard Unix permissions specified as any combination of `r w x`.

<code>setfacl -m d:u:user:permissions dir</code>	Same as above, but set a default ACL instead of an access ACL.
<code>setfacl -d -m u:user:permissions dir</code>	This applies to all commands above

<code>getfacl file</code>	Display the access (and default, if any) ACL for a file
<code>getfacl file1 setfacl --set-file=- file2</code>	Copy the ACL of <i>file1</i> and apply it to <i>file2</i>
<code>getfacl --access dir setfacl -d -M- dir</code>	Copy the access ACL of a directory and set it as default ACL

<code>chacl options</code>	Change an ACL. This command exists to provide compatibility with IRIX
----------------------------	--

<code>man acl</code>	Show the manpage about ACLs
----------------------	-----------------------------

An **inode** is a structure containing all file metadata: file type, permissions, owner, group, size, number of links, attributes, access/change/modification/deletion times, ACLs, and address where the actual file content (data) is stored. An inode does not contain the name of the file; this information is stored in the directory where the file is located (i.e. referenced). A directory contains a list of mappings between filenames and inodes.

In Linux, two types of links exist: **hard links** and **symbolic links** (aka **symlinks** or **soft links**).

The **link count** of a file is the total number of hard links to that file (i.e. to that file's inode). By default, files have a link count of 1, and directories have a link count of 2 (the directory itself, and the `.` link inside the directory). The link count of a directory is increased by one for each subdirectory (because of the `..` parent link inside the subdirectory). Once a file has no hard links pointing to it, the file is deleted, provided that no process holds the file open for reading.

	Hard link	Symbolic link
Definition	A link to an already existing inode	A path to a filename; a shortcut
Command to create it	<code>ln file hardlink</code>	<code>ln -s file symlink</code>
Link is still valid if the original file is moved or deleted	Yes (because the link still references the inode to which the original file pointed)	No (because the path now references a non-existent file)
Can link to a file in another filesystem	No (because inode numbers make sense only within a determinate filesystem)	Yes
Can link to a directory	No	Yes
Link permissions	Reflect the original file's permissions, even when these are changed	<code>rxwxrwxrwx</code>
Link attributes	- (regular file)	<code>l</code> (symbolic link)
Inode number	The same as the original file	A different inode number (since it's a different file)

`ls -li` Show a listing of the directory with the inode number for each file

`ls -l` Show a listing of the directory with the link count for each file

`df -li` Report filesystem inode usage

`find / -inum n` Find all files linked to the same inode *n*

`find / -samefile file` Find all files linked to the same inode as *file*

<code>find /path -name "foo*"</code>	Find all files and dirs, in the directory tree rooted at <i>/path</i> , whose name starts with "foo"
<code>find /path -name "foo*" -print</code>	
<code>find / -name "foo*" -exec chmod 700 {} \;</code>	Find all files and dirs whose name start with "foo" and apply permission 700 to all of them
<code>find / -name "foo*" -ok chmod 700 {} \;</code>	Find all files and dirs whose name start with "foo" and apply permission 700 to all of them, asking for confirmation
<code>find / -size +128M</code>	Find all files larger than 128 Mb
<code>find / -type f -ctime +10</code>	Find all files last changed more than 10 days ago
<code>find / -type f -perm -4000</code>	Find all files with SUID set (a possible security risk, because a shell with SUID root is a backdoor)
<code>find / -type f -newermt "May 4 2:55" -delete</code>	Find and delete all files newer than the specified timestamp. Using <code>-delete</code> is preferable to using <code>-exec rm {} \;</code>
<code>find . -type f -print -exec cat {} \;</code>	Print all files, in the current directory and under, prepending them with a filename header
<code>find . \! -name "*.gz" -type f -exec gzip {} \;</code>	Find all files, in the current directory and under, which do not have the <code>gz</code> extension, and compress them
<code>find / -xdev -type f -size +100M -exec ls -lah {} \;</code>	Find all files larger than 100 Mb in the current filesystem only and display detailed information about them
 <code>locate file</code> <code>slocate file</code>	 Locate <i>file</i> by searching the file index <code>/etc/updatedb.conf</code> , not by actually walking the filesystem. The search is fast but will only held results relative to the last rebuild of the file index
<code>updatedb</code>	Rebuild the file index
 <code>which command</code> <code>which -a command</code>	 Locate a binary executable command within the PATH Locate all matches of a command, not only the first one
 <code>whereis command</code> <code>whereis -b command</code> <code>whereis -s command</code> <code>whereis -m command</code>	 Locate the binary, source, and manpage files for a command Locate the binary files for a command Locate the source files for a command Locate the manpage files for a command
 <code>command -v command</code>	 Show the full path or alias for <i>command</i> . If no path is shown, <i>command</i> is a shell built-in (i.e. an internal feature of the shell)
 <code>compgen</code> <code>compgen -c grep '^qemu-'</code>	 Generate possible completion matches for aliases, shell built-ins, commands, directories, files, groups, user names, etc. Print a list of all QEMU commands installed and available on the machine
 <code>type command</code> <code>whence command</code> <code>where command</code>	 (Bash) (KornShell , Z shell) (Z shell) Determine whether a command is a program, a built-in, an alias, or a function
 <code>fzf</code> <code>sk</code> <code>fzy</code>	 General purpose CLI fuzzy finder Fuzzy finder Fuzzy file finder for the terminal

<code>history</code>	Show the history of command lines executed up to this moment. Commands prepended by a space will be executed but will not show up in the history. After the user logs out from Bash, history is saved into <code>~/.bash_history</code>
<code>!<i>n</i></code>	Execute command number <i>n</i> in the command line history
<code>!!</code>	Execute again the last executed command
<code>history -c</code>	Clear the command line history
<code>history -d <i>n</i></code>	Delete command number <i>n</i> from the command line history
<code>export HISTSIZE=<i>m</i></code>	Set the command line history to contain only the <i>m</i> past commands
<code>HISTTIMEFORMAT="%d/%m/%y %T "</code>	Set the command line history to also show timestamps of past commands
<code>fc</code>	Edit and execute again the last executed command
<code>fc -l -<i>n</i></code>	List the last <i>n</i> commands in the command line history
<code>alias ls='ls -lap'</code>	Set up an alias for the <code>ls</code> command
<code>alias</code>	Show all defined aliases
<code>unalias ls</code>	Remove the alias for the <code>ls</code> command
<code>unalias -a</code>	Remove all defined aliases
<code>\ls</code>	Run the non-aliased version of the <code>ls</code> command
<code>/bin/ls</code>	
<code>rm ./-rf</code>	Delete a file called "-rf". By specifying the file path, either relative or absolute, it is possible to operate on filenames which begin with a special character

Almost all Linux commands accept the option `-v` (verbose), and some commands also accept the options `-vv` or `-vvv` (increasing levels of verbosity).

All Bash built-in commands, and all commands that respect the POSIX requirements, accept the option `--` which marks in the arguments the end of options and the start of operands:

<code>grep -- -i file</code>	Search for the string "-i" in <i>file</i>
<code>rm -- -rf</code>	Delete a file called "-rf"
<code>info -- cat --</code>	Display the Info entry for <code>--</code>

A commonly used timestamp format is the **Unix epoch time**, which is defined as the number of seconds elapsed since 00:00:00 UTC (Coordinated Universal Time) 1 January 1970.

Shells		
<code>sh</code>	Thompson Shell	The first Unix shell, now obsolete. Simple command interpreter, not designed for scripting.
<code>sh</code>	Bourne Shell	Default shell for Version 7 Unix, now obsolete. On current Linux systems, <code>/bin/sh</code> is a symlink to the default shell.
<code>bash</code>	Bash (Bourne Again Shell)	GNU replacement for the Bourne Shell. Default shell for most Linux distributions.
<code>csh</code>	C shell	Shell originally designed for BSD.
<code>tcsh</code>	tcsh	C shell with additional features.
<code>ksh</code>	KornShell	Shell based on the Bourne Shell, with improvements.
<code>zsh</code>	Z shell	Shell based on the Bourne Shell, with improvements.
<code>ash</code>	Almquist shell	Lightweight shell, was the default shell for BSD.
<code>dash</code>	Debian Almquist shell	Port of Almquist shell for Debian.
<code>fish</code>	Friendly interactive shell	Interactive, user-friendly shell.

The scope of **variables** is the current shell only, while **environment variables** are visible within the current shell as well as within all subshells and Bash child processes spawned by the shell.

Environment variables are set in `/etc/environment` in the form `variable=value`.

Conventionally, variable names are lowercase while environment variable names are uppercase.

<code>set</code>	Display all set variables
<code>env</code>	Display all environment variables
<code>readonly -p</code>	Display all variables that are read-only
 <code>VAR=value</code> <code>((VAR=value))</code> <code>let "VAR=value"</code>	 Set the value of a variable. No space is allowed around the <code>=</code> sign. Spaces are allowed around <code>((</code> and <code>)</code>
<code>readonly VAR=value</code>	Set a variable as read-only, making its value unchangeable
<code>set \${VAR:=value}</code> <code>VAR=\${VAR:-value}</code>	Set a variable only if it is not already set (i.e. does not exist) or is null
<code>unset VAR</code>	Unset (i.e. delete) a variable
<code>export VAR</code>	Export a variable, making it an environment variable
<code>export VAR=value</code>	Set the value of a variable and export it
 <code>envsubst < file</code>	 Read a text file mentioning environment variables in the form <code>\$VAR</code> and output it, replacing each variable name with its value
 <code>command \$VAR</code> <code>command \${VAR}HELLO</code> <code>command "\${VAR}"</code>	 Pass a variable as argument to <i>command</i> . If other characters follow the variable name, it is necessary to specify the boundaries of the variable name via <code>{}</code> to make it unambiguous. It is recommended to double quote the variable when referencing it, to prevent interpretation of special characters (except <code>\</code> <code>\$</code> <code>`</code>) and word splitting (if the variable value contains whitespaces), which will cause unintended results
 <code>VAR=value command</code>	 Set the environment variable <i>VAR</i> to <i>value</i> for this single invocation of <i>command</i> . Note that the shell parses a line by expanding a variable before its assignment takes effect, so the following command will not output anything because <code>\$FOO</code> is still empty at the moment it is expanded: <code>FOO="hello world" echo \$FOO</code>
 <code>VAR=`command`</code> <code>VAR=\$(command)</code>	 Command substitution. Assigns to a variable the standard output resulting from <i>command</i> (which is executed in a subshell)
 <code>for i in /path/*</code> <code>do</code> <code> echo "Filename: \$i"</code> <code>done</code>	 Loop and operate through all the output tokens (in this case, echo all files in the <i>path</i>). Note that the construct <code>for i in \$(ls /path/)</code> is incorrect, because filenames containing particular characters (whitespaces, glob characters, hyphens, etc.) will cause unintended results

```
ARRAY=(val1 val2 val3)
```

Set an array (i.e. a variable containing multiple values), assigning the first three elements

```
ARRAY[3]=val4
```

Assign a fourth element in the previous array

```
echo ${ARRAY[n]}
```

Print the array element number *n*

```
echo ${ARRAY[*]}
```

Print all array elements

```
echo ${VAR:-message}
```

If variable exists and is not null, print its value, otherwise print *message*

```
echo ${VAR:+message}
```

If variable exists and is not null, print *message*, otherwise print nothing

```
VAR=$((5 + 37))
```

Evaluate a numeric expression, assigning the result to another variable

```
VAR=$((5 + 37))
```

```
VAR=$((OTHERVAR - 3))
```

```
VAR=$((OTHERVAR - 3))
```

```
VAR=`expr $OTHERVAR - 3`
```

```
[ $((VAR % 2)) -eq 0 ] && command
```

Evaluate a numeric expression, running *command* if the variable is even

```
((VAR++))
```

Increase a numeric variable by 1

```
((++VAR))
```

```
((VAR+=1))
```

```
((VAR=VAR+1))
```

```
TOKENS=($STRING)
```

String tokenizer. Splits a string stored in the variable *STRING* into tokens, according to the content of the shell variable *\$IFS*, and stores them as elements in the array *TOKENS*

```
echo ${#STRING}
```

Print a string's length (number of characters)

```
echo ${VAR,,}
```

Print a string variable in lowercase

Bash built-in variables	
\$0	Script name
\$n	<i>n</i> th argument passed to the script or function
\$@	All arguments passed to the script or function; each argument is a separate word
\$*	All arguments passed to the script or function, as a single word
\$#	Number of arguments passed to the script or function
\$?	Exit status of the last recently executed command
\${PIPESTATUS[n]}	Exit status of the <i>n</i> th command in the executed pipeline
\$\$	PID of the script in which this variable is called
\$_	PID of the last recently executed background command
\$SHLVL	Deepness level of current shell, starting with 1
\$IFS	Internal Field Separator; defines the token separators for strings, to perform word splitting after expansion. By default it has the value "space, tab, newline"
\$RANDOM	Pseudorandom integer value between 0 and 32767

Bash shell event	Files run
When a login shell is launched	/etc/profile /etc/profile.d/*.sh ~/.bash_profile ~/.bash_login ~/.profile The shell executes the system-wide profile files, then the first of the 3 user files that exists and is readable
When a login shell exits	~/.bash_logout
When a non-login shell is launched	/etc/bash.bashrc /etc/bashrc ~/.bashrc

<code>set -option</code> <code>set -o longoption</code>	Enable a Bash option
<code>set +option</code> <code>set +o longoption</code>	Disable a Bash option
<code>set -o</code>	Show the status of all Bash options
<code>set -v</code> <code>set -o verbose</code>	Print shell input lines as they are read
<code>set -x</code> <code>set -o xtrace</code>	Print command traces before execution of each command (debug mode)
<code>set -e</code> <code>set -o errexit</code>	Exit the script immediately if a command fails. Recommended option
<code>set -u</code> <code>set -o nounset</code>	Treat expansion of unset variables as an error. This avoids unintended results

There are three ways to run a script with a specific Bash option enabled:

- Run the script with `bash -option script.sh`
- Specify the shebang line in the script as `#!/bin/bash -option`
- Add the command `set -option` at the beginning of the script

<code>shopt</code>	Display the list of all shell options with their current value (on or off)
<code>shopt -s shelloption</code>	Set (enable) a specific shell option
<code>shopt -u shelloption</code>	Unset (disable) a specific shell option

Bash shell scripts must start with the shebang line `#!/bin/bash` indicating the location of the script interpreter.

Script execution	
<code>source script.sh</code> <code>. script.sh</code>	Script execution takes place in the same shell. Variables defined and exported in the script are seen by the shell when the script exits
<code>bash script.sh</code> <code>./script.sh</code> (file must be executable)	Script execution spawns a new shell

<code>command &</code>	Execute <i>command</i> in the background
<code>command1; command2</code>	Execute <i>command 1</i> and then <i>command 2</i>
<code>command1 && command2</code>	Execute <i>command 2</i> only if <i>command 1</i> executed successfully (exit status = 0)
<code>command1 command2</code>	Execute <i>command 2</i> only if <i>command 1</i> did not execute successfully (exit status > 0)
<code>(command1 && command2)</code>	Group commands together for evaluation priority
<code>(command)</code>	Run <i>command</i> in a subshell. This is used to isolate <i>command</i> 's effects, as variable assignments and other changes to the shell environment operated by <i>command</i> will not persist after <i>command</i> completes
<code>exit</code>	Terminate a script
<code>exit n</code>	Terminate a script with the specified exit status number <i>n</i> . By convention, a 0 exit status is used if the script executed successfully, a non-zero value otherwise
<code>command exit 1</code>	(To be used inside a script.) Exit the script if <i>command</i> fails
<code>/bin/true</code>	Do nothing and return immediately a status code of 0 (indicating success)
<code>/bin/false</code>	Do nothing and return immediately a status code of 1 (indicating failure)
<code>:</code>	Do nothing and return a status code of 0. Shell builtin
<code>if command</code> <code>then echo "Success"</code> <code>else echo "Failure"</code> <code>fi</code>	Run a command, then evaluate whether it exited successfully or failed
<code>function fname { commands }</code> <code>fname() { commands }</code>	Define a function. A function must be defined before it can be used in a Bash script. Argument number <i>n</i> is accessed in the body of the function via <code>\$n</code> . An advantage of functions over aliases is that functions can be passed arguments
<code>fname arg1 arg2 ...</code>	Call a function
<code>readonly -f fname</code>	Mark an already defined function as read-only, preventing it to be redefined
<code>typeset -f</code>	Show functions defined in the current Bash session
<code>readonly -p -f</code>	Show functions which are read-only
<code>shellcheck</code>	Script analyzer and debugger
<code>dialog</code>	Display shell script (terminal) dialogs for user messages and input
<code>zenity</code>	Display GTK+ graphical dialogs for user messages and input

getopts

Parse positional parameters in a shell script

getopts syntax	
while getopts abc:d: OPT do case \$OPT in a) command_a exit 0 ;; b) command_b exit 0 ;; c) command_c \$OPTARG exit 0 ;; d) command_d \$OPTARG exit 0 ;; *) default_command exit 1 ;; esac done	Definition of accepted options Matches option -a. Executes a command Matches option -c <i>argument</i>. Executes a command with argument Command to execute if none of above options applies

```
cat /etc/debian_version (Debian)
cat /etc/fedora-release (Fedora)
cat /etc/redhat-release (RHEL)
cat /etc/os-release
cat /etc/lsb-release
lsb_release -a
inxi -Sxxx
hostnamectl
```

Display Linux distribution name and version

```
screenfetch
neofetch
fastfetch
afetch
hfetch
pfetch
ufetch
neowofetch
hyfetch
nerdfetch
nextfetch
pasfetch
```

Display an ASCII art logo of the Linux distro as well as detailed system information about Linux distribution name and version, kernel version, CPU, RAM, Desktop Environment, Window Manager, screen resolution, etc.

```
rpm -qi basesystem | grep "Install Date" (Red Hat)
stat / | grep "Birth"
```

Display when, approximately, a system was installed

```
sos
```

Tool for collecting system logs and other information for debugging

```
sos report
sosreport
```

Generate an archive of diagnostic and configuration data for technical support

```
citellus.py sosreportfile
```

Perform some automated checks for troubleshooting a system

<code>watch command</code>	Execute <i>command</i> every 2 seconds
<code>watch -d -n 1 command</code>	Execute <i>command</i> every second, highlighting the differences in the output
<code>time command</code>	Execute <i>command</i> and, at its completion, write to stderr timing statistics about the run: elapsed real time between invocation and termination, user CPU time, system CPU time
<code>timeout 30s command</code>	Execute <i>command</i> and kill it after 30 seconds
<code>command ts</code>	Prepend a timestamp to each line of the output of <i>command</i>
<code>rlwrap command</code>	Readline wrapper. Executes <i>command</i> , intercepting user input to provide line editing, history, and completion
<code>run-parts directory</code>	Run all the executable files present in <i>directory</i> . Used to run the cron scripts in <code>/etc/cron.hourly/</code> , <code>/etc/cron.daily/</code> , etc.
<code>xargs command</code>	Call <i>command</i> multiple times, one for each argument found on stdin
<code>ls foo* xargs cat</code>	Print via <code>cat</code> the content of every file whose name starts by "foo"
<code>parallel command</code>	Run <i>command</i> in parallel. This is used to operate on multiple inputs, similarly to <code>xargs</code>

<code>sleep 5</code>	Pause for 5 seconds
<code>sleep \$[((\$RANDOM % 60) + 1)]s</code>	Sleep for a random time between 1 and 60 seconds
<code>sleep infinity</code>	Pause forever
<code>usleep 5000</code>	Pause for 5000 microseconds
<code>yes</code>	Output endlessly the string "y"
<code>yes string</code>	Output endlessly <i>string</i>
<code>yes fsck /dev/sda</code>	Automatically answer yes every time <code>fsck</code> asks for confirmation before fixing errors
<code>script file</code>	Generate a typescript of a terminal session. Forks a subshell and starts recording on <i>file</i> everything that is printed on terminal; the typescript ends when the user exits the subshell
<code>expect</code>	Dialogue with interactive programs according to a script, analyzing what can be expected from the interactive program and replying accordingly
<code>yarn</code>	Scenario testing tool for command line programs
<code>cmdtest</code>	Tool for black box testing of Linux command line programs
<code>busybox</code>	BusyBox, "the Swiss Army knife of Embedded Linux", an optimized multi-call binary which contains many Linux commands and utilities. Useful for system recovery if Bash built-ins or common commands have become unusable or have been removed from the system
<code>busybox applet arguments</code>	Execute <i>applet</i> , which operates as the homonym Linux command

```
test "$MYVAR" operator "value" && command
[ "$MYVAR" operator "value" ] && command
if [ "$MYVAR" operator "value" ]; then command; fi
```

Perform a test; if it results true, *command* is executed

Conditional operators			
Integer operators		File operators	
-eq value	Is equal to	-e file or -a file	Exists
-ne value	Is not equal to	-f file	Is a regular file
-lt value	Is less than	-d file	Is a directory
-le value	Is less than or equal to	-b file	Is a block special file
-gt value	Is greater than	-c file	Is a character special file
-ge value	Is greater than or equal to	-r file	Is readable
Numeric operators		-w file	Is writable
= value	Is equal to	-x file	Is executable
!= value	Is not equal to	-k file	Is sticky
< value	Is less than	-u file	Is SUID
<= value	Is less than or equal to	-g file	Is SGID
> value	Is greater than	-O file	Is owned by the Effective UID
>= value	Is greater than or equal to	-G file	Is owned by the Effective GID
Expression operators		-p file	Is a named pipe (aka FIFO)
expression1 -a expression2	Logical AND	-S file	Is a socket
expression1 -o expression2	Logical OR	-h file or -L file	Is a symbolic link
! expression	Logical NOT	-s file	Is non-zero length
\(expression \)	Priority	-N file	Was modified since last read
String operators		file1 -nt file2	Is newer than
-z	Is zero length	file1 -ot file2	Is older than
-n or nothing	Is non-zero length	file1 -ef file2	Refers to same device and inode as
= string or == string	Is equal to		
!= string	Is not equal to		
< string	Is alphabetically before		
> string	Is alphabetically after		
substr string pos len	Substring		
index string chars	Index of any chars in string		
length string	String length		
string : regex	String matches regex		
or			
match string regex			

```
expr "$MYVAR" = "39 + 3"
```

Evaluate an expression (in this case, assigns the value 42 to the variable)

```
expr string : regex
```

Return the length of the substring matching the regex

```
expr string : \(regex\)
```

Return the substring matching the regex

Operators			
Mathematical operators		Logical operators	
+	Addition	!	Logical negation
-	Subtraction	&&	Logical AND
*	Multiplication		Logical OR
/	Division	Bitwise operators	
%	Remainder	~	Bitwise negation
**	Exponentiation	&	Bitwise AND
<code>++variable</code>	Pre-increment		Bitwise OR
<code>variable++</code>	Post-increment	^	Bitwise XOR
<code>--variable</code>	Pre-decrement	<<	Left bitwise shift
<code>variable--</code>	Post-decrement	>>	Right bitwise shift
Assignment operators			
=	Assignment		
<code>operation=</code>	Operation and assignment		

Tests	
<pre>if [test 1] then [command block 1] elif [test 2] then [command block 2] else [command block 3] fi</pre>	<pre>case \$STRING in pattern1) [command block 1] ;; pattern2) [command block 2] ;; *) [command block default] ;; esac</pre>

Loops		
<pre>while [test] do [command block] done</pre> The <i>command block</i> executes as long as <i>test</i> is true	<pre>until [test] do [command block] done</pre> The <i>command block</i> executes as long as <i>test</i> is false	<pre>for item in [list] do [command block] done</pre> The <i>command block</i> executes for each <i>item</i> in <i>list</i>
<pre>i=0 while [\$i -le 7] do echo \$i let i++ done</pre>	<pre>i=0 until [\$i -gt 7] do echo \$i let i++ done</pre>	<pre>for i in 0 1 2 3 4 5 6 7 do echo \$i done</pre> <pre>for i in {0..7} do echo \$i done</pre> <pre>start=0 end=7 for i in \$(seq \$start \$end) do echo \$i done</pre> <pre>start=0 end=7 for ((i = start; i <= end; i++)) do echo \$i done</pre>
Loop jumps		
<pre>break</pre> Exit a loop	<pre>continue</pre> Jump to the next iteration	
<pre>i=0 while true do if [\$i -gt 7]; then break; fi echo \$i let i++ done</pre>	<pre>i=-9 while [\$i -lt 7] do let i++ if [\$i -lt 0]; then continue; fi echo \$i done</pre>	

<code>vi</code>	Vi, a text editor
<code>vim</code>	Vi Improved, an advanced text editor
<code>gvim</code>	Vim with GUI
<code>vimdiff file1 file2</code>	Compare two text files in Vim
<code>pico</code>	PIne COmposer, a simple text editor derived from Pine
<code>nano</code>	Simple text editor, GNU clone of Pico
<code>rnano</code>	Restricted version of Nano: does not allow the user to access the filesystem (except for files specified as argument) or to run a command shell
<code>emacs</code>	GUI text editor
<code>gedit</code>	GUI text editor
<code>ed</code>	Line-oriented text editor
<code>more</code>	Text pager. Obsolete
<code>less</code>	Text pager
<code>most</code>	Text pager with advanced features (screen split, binary viewer, etc.)

<code>g</code>	Go to the first line in the file
<code>ng</code>	Go to line number <i>n</i>
<code>G</code>	Go to the last line in the file
<code>F</code>	Go to the end of the file, and move forward automatically as the file grows
<code>CTRL C</code>	Stop moving forward
<code>-N</code>	Show line numbers
<code>-n</code>	Don't show line numbers
<code>=</code>	Show information about the file
<code>CTRL G</code>	Show current and total line number, byte, and percentage of the file read
<code>/pattern</code>	Search <i>pattern</i> forward
<code>?pattern</code>	Search <i>pattern</i> backwards
<code>&pattern</code>	Display only lines matching <i>pattern</i>
<code>n</code>	Search next occurrences forward
<code>N</code>	Search next occurrences backwards
<code>:n</code>	When reading multiple files, go to the next file
<code>:p</code>	When reading multiple files, go to the previous file
<code>R</code>	Repaint the screen
<code>V</code>	Show version number
<code>h</code>	Help
<code>q</code>	Quit

`less +command file`

Open *file* for reading, applying *command* (see list above)

`less +F --follow-name file`

Move forward, attempting periodically to reopen *file* by name; useful to keep reading a logfile that is being rotated. Note that, by default, `less` continues to read the original input file even if it has been renamed

ESC	Go to Command mode		
i	Insert text before cursor		
I	Insert text after line	and go to Insert mode	
a	Append text after cursor		
A	Append text after line		
v	Go to Visual mode, character-wise	then use the arrow keys to select a block of text	
V	Go to Visual mode, line-wise		
d	Delete selected block	gu	Switch selected block to lowercase
Y	Copy (yank) selected block into buffer	gU	Switch selected block to uppercase
w	Move to next word	\$	Move to end of line
b	Move to beginning of word	1G	Move to line 1 i.e. beginning of file
e	Move to end of word	G	Move to end of file
O	Move to beginning of line	z	RETURN Make current line the top line of the screen
CTRL G	Show current line and column number		
ma	Mark position "a". Marks a-z are local to current file, while marks A-Z are global to a specific file		
'a	Go to mark "a". If using a global mark, it also opens the specific file		
y'a	Copy (yank) from mark "a" to current line, into the buffer		
d'a	Delete from mark "a" to current line		
P	Paste buffer after current line	yy	Copy current line
P	Paste buffer before current line	yyP	Duplicate current line
x	Delete current character	D	Delete from current character to end of line
X	Delete before current character	dd	Delete current line
7dd	Delete 7 lines. Almost any command can be prepended by a number to repeat it that number of times		
u	Undo last command. Vi can undo the last command only, Vim is able to undo several commands		
.	Repeat last text-changing command		
/string	Search for <i>string</i> forward	n	Search for next match of <i>string</i>
?string	Search for <i>string</i> backwards	N	Search for previous match of <i>string</i>
:s/s1/s2/	Replace the first occurrence of <i>s1</i> with <i>s2</i> in the current line		
:s/s1/s2/g	Replace globally every occurrence of <i>s1</i> with <i>s2</i> in the current line		
:%s/s1/s2/g	Replace globally every occurrence of <i>s1</i> with <i>s2</i> in the whole file		
:%s/s1/s2/gc	Replace globally every occurrence of <i>s1</i> with <i>s2</i> in the whole file, asking for confirmation		
:5,40s/^/#/	Add a hash character at the beginning of each line, from line 5 to 40		
!!program	Replace line with output from <i>program</i>		
:r file	Read <i>file</i> and insert it after current line		
:X	Encrypt current document. Vi will automatically prompt for the password to encrypt and decrypt		
:w file	Write to <i>file</i>		
:wq	Save changes and quit		
:x			
ZZ			
:q	Quit (fails if there are unsaved changes)	:q!	Abandon all changes and quit

vi -R file Open *file* in read-only mode

cat file | vi - Open *file* in read-only mode; this is done from the shell, by having Vi read from stdin

Option	Effect
ai	Turn on auto indentation
all	Display all options
ap	Print a line after the commands <code>d c J m :s t u</code>
aw	Automatic write on commands <code>:n ! e# ^^ :rew ^} :tag</code>
bf	Discard control characters from input
dir=tmpdir	Set <i>tmpdir</i> as directory for temporary files
eb	Precede error messages with a bell
ht=8	Set terminal tab as 8 spaces
ic	Ignore case when searching
lisp	Modify brackets for Lisp compatibility
list	Show tabs and EOL characters
set listchars=tab:>-	Show tab as > for the first char and as - for the following chars
magic	Allow pattern matching with special characters
mesg	Enable UNIX terminal messaging
nu	Show line numbers
opt	Speed up output by eliminating automatic Return
para=LlPLPPPQPbpP	Set macro to start paragraphs for { } operators
prompt	Prompt : for command input
re	Simulate smart terminal on dumb terminal
remap	Accept macros within macros
report	Show the largest size of changes on status line
ro	Make file read-only
scroll=12	Set screen size as 12 lines
shell=/bin/bash	Set shell escape to /bin/bash
showmode	Show current mode on status line
slow	Postpone display updates during inserts
sm	Show matching parentheses when typing
sw=8	Set shift width to 8 characters
tags=/usr/lib/tags	Set path for files checked for tags
term	Print terminal type
terse	Print terse messages
timeout	Eliminate 1-second time limit for macros
tl=3	Set significance of tags beyond 3 characters (0 = all)
ts=8	Set tab stops to 8 for text input
wa	Inhibit normal checks before write commands
warn	Display the warning message "No write since last change"
window=24	Set text window as 24 lines
wm=0	Set automatic wraparound 0 spaces from right margin
<pre> :set option turn on an option :set nooption turn off an option :set option ? show the current value of option </pre>	
Options can also be permanently set by including them in ~/.exrc (Vi) or ~/.vimrc (Vim)	

```
SHOW DATABASES;

USE CompanyDatabase;

SELECT DATABASE();

DROP DATABASE CompanyDatabase;
```

Show all existing databases

Select a database to use

Show which database is currently selected

Delete a database

```
SHOW TABLES;

CREATE TABLE customers (
  cusid INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
  firstname VARCHAR(32), lastname VARCHAR(32), dob DATE,
  city VARCHAR(24), zipcode VARCHAR(5));

CREATE TABLE payments (
  payid INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
  date DATE, fee INT, bill VARCHAR(128), cusid INT,
  CONSTRAINT FK1 FOREIGN KEY (cusid) REFERENCES customers(cusid));
```

Show all tables from the selected database

Create tables

```
INSERT INTO customers (firstname,lastname,dob)
VALUES ('Arthur','Dent',1959-08-01), ('Trillian','',1971-03-19);

DELETE FROM customers WHERE firstname LIKE 'Zaphod';

UPDATE customers SET city = 'London' WHERE zipcode = 'L1 42HG';

CREATE INDEX lastname_index ON customers(lastname);
ALTER TABLE customers ADD INDEX lastname_index (lastname);
```

Insert new records in a table

Delete some records in a table

Modify records in a table

Create an index for faster searches

```
DESCRIBE customers;

SHOW CREATE TABLE customers;

SHOW INDEXES FROM customers;

DROP TABLE customers;
```

Describe the columns of a table

Show the code used to create a table

Show primary key and indexes of a table

Delete a table

```
ALTER TABLE customers MODIFY city VARCHAR(32);
```

Modify a column type

```
CREATE VIEW cust_view AS
SELECT * FROM customers WHERE city != 'London';
```

Create a view. Views are used similarly to tables

```
COMMIT;

ROLLBACK;
```

Commit changes to the database

Rollback the current transaction, canceling any changes done during it

```
START TRANSACTION;
BEGIN;
```

Disable autocommit for this transaction, until a COMMIT or ROLLBACK is issued

If no database has been selected for use, tables must be referenced by *databaseName.tableName*.

```
SELECT * FROM customers;
```

Select all columns from the customers table

```
SELECT firstname, lastname FROM customers LIMIT 5;
```

Select first and last name of customers, showing 5 records only

```
SELECT firstname, lastname FROM customers LIMIT 1000,5;
SELECT firstname, lastname FROM customers OFFSET 1000 LIMIT 5;
```

Select first and last name of customers, skipping the first 1000 records and showing 5 records only

```
SELECT firstname, lastname FROM customers WHERE zipcode = 'L1 42HG';
```

Select first and last name of customers whose zip code is "L1 42HG"

```
SELECT firstname, lastname FROM customers WHERE zipcode IS NOT NULL;
```

Select first and last name of customers with an existing zip code

```
SELECT * FROM customers ORDER BY lastname, firstname;
```

Select customers in alphabetical order by last name, then first name

```
SELECT * FROM customers ORDER by zipcode DESC;
```

Select customers, sorting them by zip code in reverse order

```
SELECT firstname, lastname,
TIMESTAMPDIFF(YEAR,dob,CURRENT_DATE) AS age FROM customers;
```

Select first name, last name, and calculated age of customers

```
SELECT DISTINCT city FROM customers;
```

Show all cities, retrieving each unique output record only once

```
SELECT city, COUNT(*) FROM customers GROUP BY city;
```

Show all cities and the number of customers in each city. NULL values are not counted

```
SELECT cusid, SUM(fee) FROM payments GROUP BY cusid;
```

Show all fee payments grouped by customer ID, summed up

```
SELECT cusid, AVG(fee) FROM payments GROUP BY cusid
HAVING AVG(fee)<50;
```

Show the average of fee payments grouped by customer ID, where this average is less than 50

```
SELECT MAX(fee) FROM payments;
```

Show the highest fee in the table

```
SELECT COUNT(*) FROM customers;
```

Show how many rows are in the table

```
SELECT cusid FROM payments t1 WHERE fee =
(SELECT MAX(t2.fee) FROM payments t2 WHERE t1.cusid=t2.cusid);
```

Show the customer ID that pays the highest fee (via a subquery)

```
SELECT @maxfee:=MAX(fee) FROM payments;
SELECT cusid FROM payments t1 WHERE fee = @maxfee;
```

Show the customer ID that pays the highest fee (via a user set variable)

```
SELECT * FROM customers WHERE lastname IN (SELECT lastname
FROM customers GROUP BY lastname HAVING COUNT(lastname) > 1);
```

Show the customers which have same last name as other customers

```
SELECT cusid FROM payments WHERE fee >
ALL (SELECT fee FROM payments WHERE cusid = 4242001;
```

Show the customer IDs that pay fees higher than the highest fee paid by customer ID 4242001

```
SELECT * FROM customers WHERE firstname LIKE 'Trill%';
```

Select customers whose first name matches the expression:

% = zero or more chars

_ = a single char

```
SELECT * FROM customers WHERE firstname REGEXP '^Art.*r$';
```

Select customers whose first name matches the regex

```
SELECT firstname, lastname FROM customers WHERE zipcode = 'L1 42HG'
UNION
SELECT firstname, lastname FROM customers WHERE cusid > 4242001;
```

Select customers that satisfy any of the two requirements

```
SELECT firstname, lastname FROM customers WHERE zipcode = 'L1 42HG'
INTERSECT
SELECT firstname, lastname FROM customers WHERE cusid > 4242001;
```

Select customers that satisfy both of the two requirements

```
SELECT firstname, lastname FROM customers WHERE zipcode = 'L1 42HG'
EXCEPT
SELECT firstname, lastname FROM customers WHERE cusid > 4242001;
```

Select customers that satisfy the first requirement but not the second

SQL	MySQL	Operation
<pre>SELECT customers.name, payments.bill FROM customers, payments WHERE customers.cusid = payments.cusid; SELECT customers.name, payments.bill FROM customers NATURAL JOIN payments; SELECT customers.name, payments.bill FROM customers JOIN payments USING (cusid); SELECT customers.name, payments.bill FROM customers JOIN payments ON customers.cusid = payments.cusid;</pre>	<pre>SELECT customers.name, payments.bill FROM customers [JOIN INNER JOIN CROSS JOIN] payments ON customers.cusid = payments.cusid; SELECT customers.name, payments.bill FROM customers [JOIN INNER JOIN CROSS JOIN] payments USING (cusid);</pre>	Perform a join (aka inner join) of two tables to select data that are in a relationship
<pre>SELECT customers.name, payments.bill FROM customers CROSS JOIN payments;</pre>	<pre>SELECT customers.name, payments.bill FROM customers JOIN payments;</pre>	Perform a cross join (aka Cartesian product) of two tables
<pre>SELECT customers.name, payments.bill FROM customers LEFT JOIN payments ON customers.cusid = payments.cusid;</pre>		Perform a left join (aka left outer join) of two tables, returning records matching the join condition and also records in the left table with unmatched values in the right table
<pre>SELECT customers.name, payments.bill FROM customers RIGHT JOIN payments ON customers.cusid = payments.cusid;</pre>		Perform a right join (aka right outer join) of two tables, returning records matching the join condition and also records in the right table with unmatched values in the left table

MySQL is the most used open source RDBMS (Relational Database Management System). It runs on TCP port 3306. On RHEL 7 and later it is replaced by its fork **MariaDB**, but the names of the client and of most tools remain unchanged.

`mysqld_safe`

Start the MySQL server (`mysqld`) with safety features such as restarting the server if errors occur and logging runtime information to the error logfile. This is the recommended command

`mysql_install_db` (deprecated)
`mysqld --initialize`

Initialize the MySQL data directory, create system tables, and set up an administrative account. To be run just after installing the MySQL server

`mysql_secure_installation`
`mariadb-secure-installation`

Set password for root, remove anonymous users, disable remote root login, and remove test database. To be run just after installing the MySQL server

`mysql -u root -p`
`mysql -u root -ppassword`
`mysql -u root -p -h host -P port`
`mysql -u root -p -eNB'SHOW DATABASES'`

Login to MySQL as root and prompt for the password

Login to MySQL as root with the specified password

Login to the specified remote MySQL host and port

Run an SQL command via MySQL. Flags are:

- e Run in batch mode
- N Do not print table header
- B Do not print table decoration characters +-|

`mysqldump -u root -p --all-databases > dump.sql`
`mysqldump -u root -p db > dump.sql`
`mysqldump -u root -p --databases db1 db2 > dump.sql`
`mysqldump -u root -p db table1 table2 > dump.sql`
`mysql -u root -p < dump.sql`

`mysql -u root -p db < dump.sql`

Backup all databases to a dump file

Backup a database to a dump file

Backup multiple databases to a dump file

Backup some tables of a database to a dump file

Restore all databases from a dump file (which contains a complete dump of a MySQL server)

Restore a specific database from a dump file (which contains one database)

`mysql_upgrade -u root -p`

Check all tables in all databases for incompatibilities with the current version of MySQL

`mysqlcheck`

Perform table maintenance. Each table is locked while is being processed. Options are:

- `--check` Check table for errors (default)
- `--analyze` Analyze table
- `--optimize` Optimize table
- `--repair` Repair table; can fix almost all problems except unique keys that are not unique

`mysqlcheck --check db table`

Check the specified table of the specified database

`mysqlcheck --check --databases db1 db2`

Check the specified databases

`mysqlcheck --check --all-databases`

Check all databases

<code>mysqlslap</code>	Tool for MySQL stress tests
<code>mysqltuner.pl</code>	Review the current MySQL installation configuration for performances and stability
<code>mysqlreport</code>	Generate a user-friendly report of MySQL status values. Obsolete
<code>mytop</code>	Monitor MySQL processes and queries
<code>innotop</code>	Monitor MySQL InnoDB transactions
<pre>dbs="\$(mysql -uroot -ppassword -Bse'SHOW DATABASES;') " for db in \$dbs do [operation on \$db] done</pre>	Perform an operation on each database name

```

SELECT User, Host FROM mysql.user;

CREATE USER 'user'@'host' IDENTIFIED BY 'p4ssw0rd';

DROP USER 'user'@'host';

SET PASSWORD FOR 'user'@'host' = PASSWORD('p4ssw0rd');
SET PASSWORD FOR 'user'@'host' = '*7E684A3DF6273CD1B6DE53';

SHOW GRANTS FOR 'user'@'host';

GRANT ALL PRIVILEGES ON database.* TO 'user'@'host';

REVOKE ALL PRIVILEGES ON database.* FROM 'user'@'host';

GRANT SELECT ON *.* TO 'john'@'localhost' IDENTIFIED BY 'p4ssw0rd';
GRANT SELECT ON *.* TO 'john'@'localhost' IDENTIFIED BY PASSWORD
'*7E684A3DF6273CD1B6DE53';

FLUSH PRIVILEGES;

SELECT * FROM database.table INTO OUTFILE 'file.csv'
FIELDS TERMINATED BY ',' OPTIONALLY ENCLOSED BY '"'
LINES TERMINATED BY '\n';

SELECT GROUP_CONCAT(COLUMN_NAME ORDER BY ORDINAL_POSITION)
FROM INFORMATION_SCHEMA.COLUMNS
WHERE TABLE_SCHEMA = 'database' AND TABLE_NAME = 'table';

USE database; SOURCE dump.sql;

USE database; LOAD DATA LOCAL INFILE 'file' INTO TABLE table;

SHOW CREATE TABLE table;
SHOW CREATE VIEW view;

DO SLEEP(n);
SELECT SLEEP(n);

SET PROFILING=1;
SHOW PROFILE;

statement;
statement\g
statement\G

SELECT /*!99999 comment*/ * FROM database.table;
SELECT /*!v statement*/ * FROM database.table;

\c
\! command

TEE logfile

```

List all MySQL users

Create a MySQL user (allowed to connect from *host* only) and set their password

Delete a MySQL user

Set a password for a MySQL user. The password can be specified either in plaintext or by its hash value

Show permissions for a user

Grant permissions to a user

Revoke permissions from a user; must match the already granted permission on the same database or table

Create a MySQL user and set their grants at the same time

Reload and commit the grant tables; must be run after any GRANT command

Export a table to a CSV file (without header)

Generate a comma-separated list of columns, to be used as header for a CSV file

Restore a database from a dump file

Populate a table with data from a file (one record per line, values separated by tabs)

Print the CREATE statement that created *table* or *view*

Sleep for *n* seconds

Enable profiling

Show the profile of the last executed query, with detailed steps and timing

Send an SQL statement to the server

Display result in vertical format, showing each record in multiple rows

Insert a comment

The commented *statement* is executed only if MySQL is version *v* or higher

Cancel current input

Run a shell command

Log all I/O of the current MySQL session to the specified logfile

SHOW VARIABLES; SHOW SESSION VARIABLES; SHOW LOCAL VARIABLES; SHOW GLOBAL VARIABLES; SHOW VARIABLES LIKE '%query%'; SHOW VARIABLES LIKE 'hostname'; SELECT @@hostname;	Print session variables (affecting current connection only) Print global variables (affecting global operations on the server) Print session variables that match the given pattern Print a session variable with the given name
 SET sort_buffer_size=10000; SET SESSION sort_buffer_size=10000; SET LOCAL sort_buffer_size=10000; SET @@sort_buffer_size=10000; SET @@session.sort_buffer_size=10000; SET @@local.sort_buffer_size=10000; SET GLOBAL sort_buffer_size=10000; SET @@global.sort_buffer_size=10000;	 Set a session variable Set a global variable
 SHOW STATUS; SHOW SESSION STATUS; SHOW LOCAL STATUS; SHOW GLOBAL STATUS; SHOW STATUS LIKE '%wsrep%';	 Print session status (concerning current connection only) Print global status (concerning global operations on the server) Print session status values that match the given pattern
 SHOW WARNINGS; SHOW ERRORS;	 Print warnings, errors and notes resulting from the most recent statement in the current session that generated messages Print errors resulting from the most recent statement in the current session that generated messages
 SHOW TABLE STATUS; SHOW ENGINE INNODB STATUS; SELECT * FROM information_schema.processlist; SHOW FULL PROCESSLIST; SELECT * FROM information_schema.processlist WHERE user='you';	 Print information about all tables of the current database e.g. engine (InnoDB or MyISAM), rows, indexes, data length Print statistics concerning the InnoDB engine Print the list of threads running in the local session; if run as root, print the list of threads running on the system Print the list of threads running in the local session and all other logged in sessions
 SELECT VERSION(); SELECT CURDATE(); SELECT CURRENT_DATE; SELECT CURTIME(); SELECT CURRENT_TIME; SELECT NOW(); SELECT USER();	 Print the version of the MySQL server Print the current date Print the current time Print the current date and time Print the current user@hostname that is logged in
 INSTALL COMPONENT 'file://component_validate_password'; SHOW VARIABLES LIKE 'validate_password%'; UNINSTALL COMPONENT 'file://component_validate_password';	 Install the Validate Password component Print the current settings of the Validate Password component Uninstall the Validate Password component
 \s	 Print status information about server and current connection

```
SELECT table_schema AS "Name",
SUM(data_length+index_length)/1024/1024 AS "Size in Mb"
FROM information_schema.tables GROUP BY table_schema;
```

Display the sizes of all databases in the system (counting data and indexes)

```
SELECT table_schema AS "Name",
SUM(data_length+index_length)/1024/1024 AS "Size in Mb"
FROM information_schema.tables WHERE table_schema='database';
```

Display the size of *database*

```
SELECT table_name AS "Name",
ROUND(((data_length)/1024/1024),2) AS "Data size in Mb",
ROUND(((index_length)/1024/1024),2) AS "Index size in Mb"
FROM information_schema.TABLES WHERE table_schema='database'
ORDER BY table_name;
```

Display data and index size of all tables of *database*

```
SELECT table_name, table_rows
FROM information_schema.tables WHERE table_schema='database';
```

Print an estimate of the number of rows of each table of *database*

```
SELECT SUM(data_length+index_length)/1024/1024 AS "InnoDB Mb"
FROM information_schema.tables WHERE engine='InnoDB';
```

Display the amount of InnoDB data in all databases

```
SELECT table_name, engine
FROM information_schema.tables WHERE table_schema = 'database';
```

Print name and engine of all tables in *database*

```
SELECT table_schema, table_name, create_time, update_time
FROM information_schema.tables;
```

Print creation times and last update times of all tables

```
SELECT CONCAT('KILL ',id,';')
FROM information_schema.processlist WHERE user='user'
INTO OUTFILE '/tmp/killuser'; SOURCE /tmp/killuser;
```

Kill all connections belonging to *user*

```
SELECT COUNT(1) SlaveThreadCount
FROM information_schema.processlist WHERE user='system user';
```

Distinguish between master and slave server; returns 0 on a master, >0 on a slave

```
SELECT ROUND(SUM(CHAR_LENGTH(field)<40)*100/COUNT(*),2)
FROM table;
```

Display the percentage of rows on which the string *field* is shorter than 40 chars

```
SELECT CHAR_LENGTH(field) AS Length, COUNT(*) AS Occurrences
FROM table GROUP BY CHAR_LENGTH(field);
```

Display all different lengths of string *field* and the number of times they occur

```
SELECT MAX(CHAR_LENGTH(field)) FROM table;
```

Display the longest string stored in *field*

```
SHOW FULL TABLES IN database WHERE table_type LIKE 'VIEW';
```

Display the list of views in *database*

```
SELECT "Table 1" AS `set`, t1.* FROM table1 t1 WHERE
ROW(t1.col1, t1.col2, t1.col3) NOT IN (SELECT * FROM table2)
UNION ALL
SELECT "Table 2" AS `set`, t2.* FROM table2 t2 WHERE
ROW(t2.col1, t2.col2, t2.col3) NOT IN (SELECT * FROM table1)
```

Display the differences between the contents of two tables *table1* and *table2* (assuming the tables are composed of 3 columns each)

How to resync a master-slave replication

1. On the master, on terminal 1:

```
mysql -uroot -p
RESET MASTER;
FLUSH TABLES WITH READ LOCK;
SHOW MASTER STATUS;
```

Note the values of MASTER_LOG_FILE and MASTER_LOG_POS; these values will need to be copied on the slave
2. On the master, on terminal 2:

```
mysqldump -uroot -p --all-databases > /path/to/dump.sql
```

It is not necessary to wait until the dump completes
3. On the master, on terminal 1:

```
UNLOCK TABLES;
```
4. Transfer the dump file from the master to the slave
5. On the slave:

```
mysql -uroot -p
STOP SLAVE;
SOURCE /path/to/dump.sql;
RESET SLAVE;
CHANGE MASTER TO MASTER_LOG_FILE='mysql-bin.nnnnnn', MASTER_LOG_POS=mm;
START SLAVE;
SHOW SLAVE STATUS;
```

How to recover the MySQL root password

1. Stop the MySQL server
2. Restart the MySQL server skipping the grant tables

```
mysqld_safe --skip-grant-tables --skip-networking &
```
3. Connect to the MySQL server passwordlessly

```
mysql -uroot
```
4. Reload the grant tables

```
FLUSH PRIVILEGES;
```
5. Change the root password

```
SET PASSWORD FOR 'root'@'localhost' = PASSWORD('newpassword'); (v5)
ALTER USER 'root'@'localhost' IDENTIFIED BY 'newpassword'; (v8)
```
6. Stop the MySQL server and restart it normally

PostgreSQL (aka **Postgres**) is an open source object-relational database. By default it listens for connections on TCP port 5432.

<code>\list</code>	List all databases
<code>\l</code>	
<code>\list+</code>	List all databases, displaying database size and description
<code>\l+</code>	
<code>\connect database</code>	Connect to <i>database</i>
<code>\c database</code>	
<code>\q</code>	Quit

How to set up PostgreSQL with a database owned by a user

- | | |
|--|---|
| 1. Set up PostgreSQL | <code>postgresql-setup initdb</code> |
| 2. Change the password of the postgres shell user | <code>passwd postgres</code> |
| 3. Create the <i>user</i> shell user | <code>useradd user</code> |
| 4. Switch to the postgres shell user and connect to PostgreSQL | <code>su - postgres</code>
<code>psql -U postgres</code> |
| 5. In PostgreSQL, create the <i>user</i> | <code>CREATE ROLE user WITH LOGIN;</code>
<code>\password user</code>
<code>\q</code> |
| 6. Create a <i>database</i> owned by <i>user</i> | <code>createdb -E utf8 -l C -T template0 database -O user</code> |
| 7. Switch to the postgres shell user and connect to PostgreSQL | <code>su - postgres</code>
<code>psql -U postgres</code> |
| 8. In PostgreSQL, grant the necessary privileges on <i>database</i> to <i>user</i> | <code>GRANT ALL PRIVILEGES ON DATABASE database TO user;</code>
<code>\q</code> |
| 9. Verify that <i>user</i> can login to PostgreSQL | <code>su - user</code>
<code>psql -U user -W</code> |

The **X Window System** (aka **X11** or **X**) is a windowing system for Linux and UNIX-like OSes, providing a basic framework for GUI applications via a client-server model. The open source implementation is the **X.Org Server**.

A **display manager** provides a login screen to enter an X session and introduces the user to the **desktop environment** (e.g. GNOME, KDE, CDE, Enlightenment).

Display Manager	Configuration files		Display Manager greeting screen
xdm X Display Manager	/etc/x11/xdm/Xaccess /etc/x11/xdm/Xresources /etc/x11/xdm/Xservers /etc/x11/xdm/Xsession /etc/x11/xdm/Xsetup_0 /etc/x11/xdm/xdm-config	Control of inbound requests from remote hosts Configuration settings for X applications and the login screen Association of X displays with local X server software, or with X terminals via XDMCP Script launched by xdm after login Script launched before the graphical login screen Association of all xdm configuration files	Defined in /etc/x11/xdm/Xresources by the line: xlogin*greeting: \ Debian GNU/Linux (CLIENTHOST)
gdm GNOME Display Manager	/etc/gdm/gdm.conf or /etc/gdm/custom.conf		Configured via gdmsetup
kdm KDE Display Manager	/etc/kde/kdm/kdmrc		Configured via kdm_config

/etc/X11/xorg.conf	Configuration file for X
~/.Xresources	Configuration settings for X applications, in the form <i>program*resource: value</i>
\$DISPLAY	Environment variable defining the display name of the X server, in the form <i>hostname:displaynumber.screennumber</i>

The following line in /etc/inittab instructs `init` to launch XDM at runlevel 5:

```
x:5:respawn:/usr/X11R6/bin/xdm -nodaemon
```

The following lines in /etc/sysconfig/desktop define GNOME as the default Display Environment and Display Manager:

```
desktop="gde"
displaymanager="gdm"
```

<code>/etc/init.d/xdm start</code> <code>/etc/init.d/gdm start</code> <code>/etc/init.d/kdm start</code>	Start the appropriate Display Manager
<code>startx</code>	Initialize an X session
<code>xorgconfig</code> (Debian) <code>Xorg -configure</code> (Red Hat)	Configure X (text mode)
<code>xorgcfg</code> (Debian) <code>system-config-display</code> (Red Hat)	Configure X (graphical mode)
<code>xhost + 10.3.3.3</code> <code>xhost - 10.3.3.3</code>	Add or remove 10.3.3.3 to the list of hosts allowed making X connections to the local machine
<code>switchdesk gde</code>	Switch to the GDE Display Manager at runtime
<code>X -version</code>	Show which version of X is running
<code>xdotool</code>	X automation tool
<code>xdotool getwindowfocus</code>	Get the ID of the currently focused window (if run in command line, it is the terminal where this command is typed)
<code>xdotool selectwindow</code>	Pop up an X cursor and get the ID of the window selected by it
<code>xdotool key --window 12345678 Return</code>	Simulate a RETURN keystroke inside window ID 12345678
<code>xprop</code>	X property displayer. Pops up a cursor to select a window
<code>xprop grep WM_CLASS</code>	Get process name and GUI application name of the selected window
<code>xrandr</code>	Configuration utility for the RandR (Resize and Rotate) X extension
<code>xrandr -q</code>	Show screen(s) size and resolution
<code>xrandr --output eDP1 --right-of VGA1</code>	Extend the screen on a VGA physical monitor situated to the left
<code>xsel</code>	Manipulate the X selection (primary, secondary, and clipboard)
<code>xsel -b < file</code>	Copy the contents of a file to the X clipboard
<code>xsel -b -a < file</code>	Append the contents of a file to the X clipboard
<code>xsel -b -o</code>	Output onscreen the contents of the X clipboard
<code>xset</code>	Configuration utility for X
<code>xset r rate 200 50</code>	Set key repeat delay to 200 ms and rate to 50 characters per second
<code>xclip</code>	X clipboard tool
<code>cat file xclip -i</code>	Copy the contents of a file to the X clipboard
<code>xautolock</code>	Run a program in case of user inactivity during a X session
<code>xdpyinfo</code>	Display information about the X server
<code>xwininfo</code>	Display information about windows
<code>xosview</code>	Monitor able to display several system parameters (CPU usage, memory usage, load average, page swapping, interrupts, battery level, etc.)

<code>xmodmap</code>	Display and edit the keyboard modifier map and keymap table used by X applications
<code>loadkeys</code>	Load keyboard translation tables
<code>kbdcontrol</code>	Control and configure the keyboard
<code>xkbcomp</code> (Red Hat) <code>ckbcomp</code> (Ubuntu)	Compile a XKB keyboard description into a format understood by <code>kbdcontrol</code> and <code>loadkeys</code>
<code>mkfontdir</code>	Catalog the newly installed fonts in the new directory
<code>xset fp+ /usr/local/share/fonts</code>	Dynamically add new installed fonts in <code>/usr/local/share/fonts</code> to the X server
<code>xfs</code>	Start the X font server
<code>fc-cache -f</code>	Scan the font directory and force a rebuild of the font information cache. To be run after adding new font files to the directory <code>/usr/local/share/fonts/</code>
<code>fc-list</code>	Print the list of fonts installed on the system
<code>fc-scan fontfile.ttf</code>	Get information about a font

Main		Latin 1				Latin 2	
BackSpace	ff08	space	0020	questiondown	00bf	Aogonek	01a1
Tab	ff09	exclam	0021	Agrave	00c0	breve	01a2
Linefeed	ff0a	quotedbl	0022	Aacute	00c1	Lstroke	01a3
Clear	ff0b	numbersign	0023	Acircumflex	00c2	Lcaron	01a5
Return	ff0d	dollar	0024	Atilde	00c3	Sacute	01a6
Pause	ff13	percent	0025	Adiaeresis	00c4	Scaron	01a9
Scroll_Lock	ff14	ampersand	0026	Aring	00c5	Scedilla	01aa
Sys_Req	ff15	apostrophe	0027	AE	00c6	Tcaron	01ab
Escape	ff1b	quoteright	0027	Ccedilla	00c7	Zacute	01ac
Delete	ffff	parenleft	0028	Egrave	00c8	Zcaron	01ae
		parenright	0029	Eacute	00c9	Zabovedot	01af
		asterisk	002a	Ecircumflex	00ca	aogonek	01b1
		plus	002b	Ediaeresis	00cb	ogonek	01b2
		comma	002c	Igrave	00cc	lstroke	01b3
		minus	002d	Iacute	00cd	lcaron	01b5
		period	002e	Icircumflex	00ce	sacute	01b6
		slash	002f	Idiaeresis	00cf	caron	01b7
		0 - 9	0030 - 0039	ETH	00d0	scaron	01b9
		colon	003a	Eth	00d0	scedilla	01ba
		semicolon	003b	Ntilde	00d1	tcaron	01bb
		less	003c	Ograve	00d2	zacute	01bc
		equal	003d	Oacute	00d3	doubleacute	01bd
		greater	003e	Ocircumflex	00d4	zcaron	01be
		question	003f	Otilde	00d5	zabovedot	01bf
		at	0040	Odiaeresis	00d6	Racute	01c0
		A - Z	0041 - 005a	multiply	00d7	Abreve	01c3
		bracketleft	005b	Oslash	00d8	Lacute	01c5
		backslash	005c	Ooblique	00d8	Cacute	01c6
		bracketright	005d	Ugrave	00d9	Ccaron	01c8
		asciicircum	005e	Uacute	00da	Eogonek	01ca
		underscore	005f	Ucircumflex	00db	Ecaron	01cc
		grave	0060	Udiaeresis	00dc	Dcaron	01cf
		quoteleft	0060	Yacute	00dd	Dstroke	01d0
		a - z	0061 - 007a	THORN	00de	Nacute	01d1
		braceleft	007b	Thorn	00de	Ncaron	01d2
		bar	007c	ssharp	00df	Odoubleacute	01d5
		braceright	007d	agrave	00e0	Rcaron	01d8
		asciitilde	007e	aacute	00e1	Uring	01d9
		nobreakspace	00a0	acircumflex	00e2	Udoubleacute	01db
		exclamdown	00a1	atilde	00e3	Tcedilla	01de
		cent	00a2	adiaeresis	00e4	racute	01e0
		sterling	00a3	aring	00e5	abreve	01e3
		currency	00a4	ae	00e6	lacute	01e5
		yen	00a5	ccedilla	00e7	cacute	01e6
		brokenbar	00a6	egrave	00e8	ccaron	01e8
		section	00a7	eacute	00e9	eogonek	01ea
		diaeresis	00a8	ecircumflex	00ea	ecaron	01ec
		copyright	00a9	ediaeresis	00eb	dcaron	01ef
		ordfeminine	00aa	igrave	00ec	dstroke	01f0
		guillemotleft	00ab	iacute	00ed	nacute	01f1
		notsign	00ac	icircumflex	00ee	ncaron	01f2
		hyphen	00ad	idiaeresis	00ef	odoubleacute	01f5
		registered	00ae	eth	00f0	rcaron	01f8
		macron	00af	ntilde	00f1	uring	01f9
		degree	00b0	ograve	00f2	udoubleacute	01fb
		plusminus	00b1	oacute	00f3	tcedilla	01fe
		twosuperior	00b2	ocircumflex	00f4	abovedot	01ff
		threesuperior	00b3	otilde	00f5		
		acute	00b4	odiaeresis	00f6		
		mu	00b5	division	00f7		
		paragraph	00b6	oslash	00f8		
		periodcentered	00b7	ooblique	00f8		
		cedilla	00b8	ugrave	00f9		
		onesuperior	00b9	uacute	00fa		
		masculine	00ba	ucircumflex	00fb		
		guillemotright	00bb	udiaeresis	00fc		
		onequarter	00bc	yacute	00fd		
		onehalf	00bd	thorn	00fe		
		threequarters	00be	ydiaeresis	00ff		

This table is derived from `keysymdef.h`, which defines keysym codes (i.e. characters or functions associated with each key in the X Window System) as `XK_key` and its hex value. The `key` can be passed as argument to the `xdotool key` command.

/etc/passwd User accounts	
root:x:0:0:/root:/bin/bash bin:x:1:1:/bin:/bin/bash jdoe:x:500:100:John Doe,,555-1234,,:/home/jdoe:/bin/bash	
1	2 3 4 5 6 7
1	Login name
2	Hashed password (obsolete), or x if shadow passwords are used i.e. the password is in /etc/shadow
3	UID – User ID
4	GID – Default Group ID
5	GECOS field – Information about the user: Full name, Room number, Work phone, Home phone, Other
6	Home directory of the user
7	Login shell (if set to /sbin/nologin or /bin/false, user will be unable to log in)

/etc/shadow User passwords	
root:\$6\$qk8JmJHf\$X9GfOZ/i9LZP3K1du6.D3cx2pXA:15537:0:99999:7::: bin*:15637:0:99999:7::: jdoe:!!\$6\$YOiH1otQ\$KxeiUKHEXK8e3jCUdw9Rxy3Wu53:15580:0:99999:7::15766:	
1	2 a b c 3 4 5 6 7 8 9
1	Login name
2	Hashed password (* if account is disabled, ! or !! if no password is set, prefixed by ! if the account is locked). Composed of the following subfields separated by \$: a Hashing algorithm: 1 = MD5, 2a = Blowfish, 5 = SHA256, 6 = SHA512 (recommended), y = yescrypt b Random salt, up to 16 chars long. This is to thwart password cracking attempts based on rainbow tables c String obtained by hashing the user's plaintext password concatenated to the stored salt
3	Date of last password change (in number of days since 1 January 1970)
4	Days before password may be changed; if 0, user can change the password at any time
5	Days after which password must be changed
6	Days before password expiration that user is warned
7	Days after password expiration that account is disabled
8	Date of account disabling (in number of days since 1 January 1970)
9	Reserved field

/etc/group Group accounts	
root:x:0:root jdoe:x:501 staff:x:530:jdoe,asmith	
1	2 3 4
1	Group name
2	Encrypted password, or x if password is in /etc/gshadow
3	GID – Group ID
4	Group members (if this is not their Default Group)

/etc/gshadow Group passwords	
root::root:root jdoe:!:: staff:0cfz7IqLhW19i::root,jdoe	
1	2 3 4
1	Group name
2	Encrypted password, or ! if no password is set (default)
3	Group administrators
4	Group members

Files /etc/shadow and /etc/gshadow have permission 000 and hence are readable only by the root user.

<code>useradd -m user</code>	Create a user account, creating and populating their homedir from <code>/etc/skel</code>
<code>useradd -mc "Name Surname" user</code>	Create a user account, specifying their full name
<code>useradd -ms /bin/ksh user</code>	Create a user account, specifying their login shell
<code>useradd -D</code>	Show default values for user account creation, as specified in <code>/etc/login.defs</code> and <code>/etc/default/useradd</code>
<code>usermod -c "Name Surname" user</code>	Modify the GECOS field of a user account
<code>usermod -L user</code>	Lock a user account
<code>usermod -U user</code>	Unlock a user account
Many options for <code>usermod</code> and <code>useradd</code> are the same.	
<code>userdel -r user</code>	Delete a user and their homedir
<code>chfn user</code>	Change the GECOS field of a user
<code>chsh user</code>	Change the login shell of a user
<code>passwd user</code>	Change the password of a user
<code>passwd -l user</code>	Lock a user account
<code>passwd -S user</code>	Show information about a user account: username, account status (L=locked, P=password, NP=no password), date of last password change, min age, max age, warning period, inactivity period in days
<code>chage -E 2022-02-14 user</code>	Change the password expiration date; account will be locked at that date
<code>chage -d 13111 user</code>	Change the date (in number of days since 1 January 1970) of last password change
<code>chage -d 0 user</code>	Force the user to change password at their next login
<code>chage -M 30 user</code>	Change the max number of days during which a password is valid
<code>chage -m 7 user</code>	Change the min number of days between password changes
<code>chage -W 15 user</code>	Change the number of days before password expiration that the user will be warned
<code>chage -I 3 user</code>	Change the number of days after password expiration before the account is locked
<code>chage -l user</code>	List password aging information for a user
<code>chpasswd</code>	Tool for batch update of passwords. Reads from stdin a list of <code>username:password</code>
<code>vipw</code> <code>vigr</code>	Edit manually <code>/etc/passwd</code> , <code>/etc/shadow</code> , <code>/etc/group</code> , or <code>/etc/gshadow</code>
<code>adduser</code> <code>deluser</code>	User-friendly front-end commands for user management
<code>system-config-users</code> (Red Hat)	GUI for user and group management

<code>groupadd group</code>	Create a group
<code>groupmod -n newname oldname</code>	Change a group name
<code>groupdel group</code>	Delete a group
<code>gpasswd group</code>	Set or change the password of a group
<code>gpasswd -a user group</code>	Add a user to a group
<code>gpasswd -d user group</code>	Delete a user from a group
<code>gpasswd -A user group</code>	Add a user to the list of administrators of the group
<code>addgroup</code> <code>delgroup</code>	User-friendly front-end commands for group management

On a system, every user is identified by a numeric **UID (User ID)**, and every group by a numeric **GID (Group ID)**.
UID 0 is assigned to the superuser.

UIDs from 0 to 99 should* be reserved for static allocation by the system and not be created by applications.

UIDs from 100 to 499 should* be reserved for dynamic allocation by the superuser and post-install scripts.

UIDs for user accounts start from 500 (Red Hat) or 1000 (SUSE, Debian).

* as recommended by the Linux Standard Base core specifications

A process has an effective, saved, and real UID and GID.

Effective UID	Used for most access checks, and as the owner for files created by the process. An unprivileged process can change its effective UID only to either its saved UID or its real UID.
Saved UID	Used when a process running with elevated privileges needs to temporarily lower its privileges. The process changes its effective UID (usually root) to an unprivileged one, and its privileged effective UID is copied to the saved UID. Later, the process can resume its elevated privileges by resetting its effective UID back to the saved UID.
Real UID	Used to identify the real owner of the process and affect the permissions for sending signals. An unprivileged process can signal another process only if the sender's real or effective UID matches the receiver's real or saved UID. Child processes inherit the credentials from the parent, so they can signal each other.

<code>whoami</code>	Print your username (as effective UID)
<code>id</code>	Print your real and effective UID and GID, and the groups of which you are a member
<code>id user</code>	Print real and effective UID and GID, and group membership information, about <i>user</i>
<code>id -u</code>	Print your effective UID
<code>who</code>	Print the list of users logged into the system
<code>w</code>	Print the list of users logged into the system, and what they are doing
<code>users</code>	Print names of currently logged in users
<code>finger user@host</code>	Print information about <i>user</i> on <i>host</i> . Shows the content of the files <code>.plan</code> and <code>.project</code> in <i>user's</i> home directory. Deprecated in RHEL 7 and later
<code>pinky user@host</code>	Print information about <i>user</i> on <i>host</i> . Lightweight version of <code>finger</code>
<code>rwho</code>	Print information about currently logged in users for all hosts on the local network
<code>rusers</code>	Print names of currently logged in users for all hosts on the local network

Sudo is a mechanism that allows running a command as another user. Sudo access rights are defined in the sudoers files `/etc/sudoers` and `/etc/sudoers.d/*`; these files must be edited only via `visudo`.

Commands run by sudo users are logged via syslog on `/var/log/auth.log` (Debian) or `/var/log/secure` (Red Hat).

<code>sudo -u user command</code>	Run <i>command</i> as <i>user</i>
<code>sudo command</code>	Run <i>command</i> as root
<code>sudo -u root command</code>	
<code>sudo su -</code> <code>sudo -i</code>	Login on an interactive shell as root
<code>sudo -u user -s</code>	Login as <i>user</i> with a shell, even if the user's shell is <code>/sbin/nologin</code> or similar
<code>sudo -l</code>	List the allowed commands for the current user
<code>sudo !!</code>	Run again the last command, but this time as root
<code>sudoedit /etc/passwd</code> <code>sudo -e /etc/passwd</code>	Safely edit a file (in this case, <code>/etc/passwd</code>) according to security policies. It is recommended to have nonprivileged users run this command instead of sudoing text editors as root on protected files, because the text editor might spawn a shell, causing security issues
<code>visudo</code> <code>visudo -c</code>	Safely edit the sudoers file Check the sudoers file for syntax errors, unused aliases, etc.
<code>man sudoers</code>	Show sudoers syntax as well as examples
<code>su user</code> <code>su</code> <code>su root</code> <code>su -c "fdisk -l"</code> <code>su -</code> <code>su -l</code>	Run a shell as <i>user</i> Run a shell as root Pass a single command to the shell Spawn a login shell, hence ensuring that login scripts are run and the correct environment variables are set. Recommended option
<code>gksudo -u user command</code> <code>gksu -u user -l</code>	GUI front-ends to <code>sudo</code> and <code>su</code> , used to run an X Window application as <i>user</i> . Pops up a requester prompting for <i>user</i> 's password
<code>pkexec --user user command</code>	Run <i>command</i> as <i>user</i>
<code>runuser -u user command</code>	Run <i>command</i> as <i>user</i> . Can be launched only by root
<code>doas -u user command</code>	Run <i>command</i> as <i>user</i> . Simpler version of <code>sudo</code> ; access rights are configured in the file <code>/etc/doas.conf</code>

`chvt n``CTRL ALT Fn`Make `/dev/tty $n$`  the foreground terminal`clear``CTRL L`

Clear the terminal screen

`vlock
away`

Lock the virtual console (terminal)

`tty`Print your terminal device (e.g. `/dev/tty1`, `/dev/pts/1`)`stty`

Change or display terminal line settings

`stty -ixon`

Disable XON/XOFF flow control

`tput`

Query a terminfo i.e. a database storing terminal capabilities

`tput cols`

Output the number of columns of the current terminal window

`tput lines`

Output the number of lines of the current terminal window

`tset``reset``tput init`

Initialize the terminal

<code>screen</code>	Screen manager that multiplexes a single virtual VT100/ANSI terminal between multiple processes or shells. Normally, when the connection to a terminal is lost (e.g. because the terminal is closed manually, the user logs out, or the remote SSH session goes into timeout), a SIGHUP is sent to the shell and from there to all running child processes which are therefore terminated. Instead, the <code>screen</code> command starts an interactive shell screen session, to which the user will be able to reattach later
<code>screen -S sessionname</code>	Start a screen session with the specified session name
<code>screen command</code>	Start <i>command</i> in a screen session; session will end when <i>command</i> exits
<code>screen -list</code>	Show the list of detached screen sessions
<code>screen -r pid.tty.host</code>	Resume a detached screen session
<code>screen -r owner/pid.tty.host</code>	
<code>screen -R</code>	Resume the last detached screen session
<code>screen -d -R sessionname</code>	Detach a remote screen session and reattach your current terminal to it
CTRL A	Send a command to the window manager: 0 ... 9 Switch between screen sessions c Create a new screen session ? Show help
<code>tmux</code>	Terminal multiplexer. Creates a new session
<code>tmux attach-session</code>	Attach to last session
<code>reptyr</code>	Attach an existing running program to a new terminal

How to detach an already running job that was not started in a screen session

(this procedure detaches the job from its parent shell, so that the job will not be killed when the terminal is closed)

1. **CTRL Z** Suspend the job
 2. `bg` Send the job to background
 3. `jobs` Show the number (let us assume is *n*) of the backgrounded job
 4. `disown -h %n` Mark job *n* so it will not receive a SIGHUP from its parent shell
- or
1. `screen` Start a new screen session
 2. `reptyr pid` Attach the job with process ID *pid* to the new terminal (screen session)

<code>write user</code>	Write interactively a message to the terminal of <i>user</i> (which must be logged in)
<code>echo "Message" write user</code>	Write a message to the terminal of <i>user</i> (which must be logged in)
<code>wall</code>	Write interactively a message to the terminal of all logged in users
<code>echo "Message" wall</code>	Write a message to the terminal of all logged in users
<code>talk user</code>	Open an interactive chat session with <i>user</i> (which must be logged in)
<code>mesg</code>	Display your current message permission status. You can allow or disallow other users to send you messages by setting or removing the group write permission of your terminal device, which is owned by system group <code>tty</code> . The root user is always able to message users, regardless of their permission status
<code>mesg y</code> <code>chmod g+w \$(tty)</code>	Allow the other users to message you via <code>write</code> , <code>wall</code> , and <code>talk</code>
<code>mesg n</code> <code>chmod g-w \$(tty)</code>	Disallow the other users to message you via <code>write</code> , <code>wall</code> , and <code>talk</code>

cron is a job scheduler, allowing repeated execution of commands specified in crontab files. The `crond` daemon checks the crontab files every minute and runs the command as the specified user at the specified times. It is not necessary to restart `crond` after the modification of a crontab file, as the changes will be reloaded automatically.

If `/etc/cron.allow` exists, only users listed therein can access the service.
If `/etc/cron.deny` exists, all users except those listed therein can access the service.
If none of these files exist, all users can access the service.

<code>/etc/crontab</code> <code>/etc/cron.d/*</code>	System-wide crontab files
<code>/etc/cron.hourly/</code> <code>/etc/cron.daily/</code> <code>/etc/cron.weekly/</code> <code>/etc/cron.monthly/</code>	Scripts placed in these directories will be automatically executed on the specified periods
<code>/var/spool/cron/user</code>	Crontab of <i>user</i> . This file has the same format as the system-wide crontab files, except that the "user" field is not present
<code>crontab -e</code>	Edit your user crontab file
<code>crontab -l</code>	List the contents of your crontab file
<code>crontab -e -u user</code>	Edit the crontab file of another <i>user</i> (command available only to the superuser)

/etc/crontab						
#	m	h	dom	mon	dow	user command
25	6	*	*	1	root	/opt/script1.sh
every Monday at 6:25 AM						
*/5	16	*	*	*	root	/opt/script2.sh
from 4:00 to 4:55 PM every 5 minutes every day						
0,30	7	25	12	*	jdoe	/home/jdoe/foo.sh
at 7:00 and 7:30 AM on 25 th December						
3	17	*	*	1-5	root	/bin/rm /tmp/abc.o
at 5:03 PM every day, from Monday to Friday						
m	minutes					
h	hours					
dom	day of month (1-31)					
mon	month (1-12 or jan-dec)					
dow	day of week (0-7 or sun-sat; 0=7=Sunday)					
user	User as whom the command will be executed					
command	Command that will be executed at the specified times					

The `crond` daemon also runs **anacron** jobs, which allow execution of periodic jobs on a machine that is not always powered on, such as a laptop. Only the superuser can schedule anacron jobs, which have a granularity of one day (vs one minute for cron jobs).

`/var/spool/anacron/jobid` Date of the last execution of the anacron job identified by *jobid*

/etc/anacrontab				
#	period	delay	job-identifier	command
7	10	cron.weekly	/opt/script3.sh	If the job has not been run in the last 7 days, wait 10 minutes and then execute the command
period	period, in days, during which the command was not executed			
delay	delay to wait, in minutes, before execution of the command			
job-identifier	job identifier in anacron messages; should be unique for each anacron job			
command	command that will be executed			

at is used for scheduled execution of commands that must run only once. These commands are executed by the `atd` daemon.

If `/etc/at.allow` exists, only users listed therein can access the service.

If `/etc/at.deny` exists, all users except those listed therein can access the service.

If none of these files exist, no user except the superuser can access the service.

```
at 2:00pm tomorrow script.sh
at -f listofcommands.txt 5:00pm
echo "systemctl start firewalld" | at now + 2 minutes
```

Execute a command once at the specified time
(absolute or relative)

```
atq
at -l
```

List the scheduled jobs

```
atrm 3
at -d 3
```

Remove job number 3 from the list

<code>systemctl list-timers</code>	List all active Systemd timers
<code>systemctl list-timers --all</code>	List all loaded timers, active and inactive
<code>/etc/systemd/system/</code>	Directory containing timers and associated services
<code>batch</code>	Schedule execution of a command for when the system is not too charged. Reads a command from stdin and runs it when the system load average falls below 0.8
<code>fcron</code>	Daemon similar to cron that executes scheduled tasks
<code>jobber</code>	Alternative to cron featuring job execution history and advanced error handling

<code>cc</code>	C compiler								
<code>gcc</code>	GNU C and C++ compiler								
<code>g++</code>	GNU C++ compiler								
<code>clang</code>	LLVM Clang C and Objective-C compiler								
<code>llc</code>	LLVM source inputs compiler into assembly language								
<code>ld</code>	GNU linker. Generates an executable file from object files created during compilation								
<code>gdb</code>	GNU Debugger. Displays what is happening inside a program while it executes. It must be used on programs compiled with a compiler belonging to the GNU toolchain								
<code>dbx</code>	Source-level debugger								
<code>lldb</code>	LLVM debugger. It must be used on programs compiled with a compiler belonging to the LLVM Project toolchain								
<code>ddd</code>	Data Display Debugger. GUI for various debuggers such as GDB, DBX, etc.								
<code>nasm</code>	Netwide Assembler, a portable 80x86 assembler								
<code>yasm</code>	Yasm Modular Assembler, a rewrite of <code>nasm</code> under the BSD License								
<code>gasp</code>	GNU Assembler Preprocessor, a preprocessor for assembly programs								
<code>make</code>	Utility for automatic compiling, re-compiling, and installation of multi-file programs. It determines automatically which parts of a multi-file program need recompiling. The relationships about these parts, and the commands that must be used to update them, are described in a Makefile (<code>./Makefile</code> by default). The Makefile also describes targets that are going to be used as arguments to the <code>make</code> command to perform the desired action, e.g.: <table><tr><td><code>all</code></td><td>Compile the whole program</td></tr><tr><td><code>install</code></td><td>Compile the program and install it, copying the executable file and all accessory files (libraries, manual, etc.) in their final destination directory for actual use</td></tr><tr><td><code>clean</code></td><td>Delete all temporary files in the current directory that are normally created by the compilation of the program, without deleting configuration files</td></tr><tr><td><code>dist</code></td><td>Create a distribution tarfile for the program</td></tr></table>	<code>all</code>	Compile the whole program	<code>install</code>	Compile the program and install it, copying the executable file and all accessory files (libraries, manual, etc.) in their final destination directory for actual use	<code>clean</code>	Delete all temporary files in the current directory that are normally created by the compilation of the program, without deleting configuration files	<code>dist</code>	Create a distribution tarfile for the program
<code>all</code>	Compile the whole program								
<code>install</code>	Compile the program and install it, copying the executable file and all accessory files (libraries, manual, etc.) in their final destination directory for actual use								
<code>clean</code>	Delete all temporary files in the current directory that are normally created by the compilation of the program, without deleting configuration files								
<code>dist</code>	Create a distribution tarfile for the program								
<code>shc</code>	Shell script compiler, used to prevent a shell script from inspection or modification. It encrypts a shell script, generates C source code, and compiles the C code into a stripped binary executable file								
<code>valgrind</code>	Dynamic analysis tool for memory management bug detection, memory leak detection, and software profiling								
<code>patch</code>	Apply or remove a patch								
<code>lsdiff</code>	List the files which are modified in a patch								
<code>objdump</code>	Display information from an object file								
<code>readelf</code>	Display information about an ELF file								

<code>bc</code>	Calculator featuring arbitrary precision arithmetic
<code>dc</code>	Calculator featuring arbitrary precision arithmetic, in Reverse Polish Notation
<code>factor</code>	Find the prime factors of an integer
<code>units</code>	Convert quantities between different units
<code>seq</code>	Print a sequence of numbers
<code>seq -s* <i>n</i> bc</code>	Calculate the factorial of <i>n</i>
<code>jot</code>	Print increasing, decreasing, random, or redundant numbers or other data
<code>jot <i>n a b</i></code>	Print <i>n</i> numbers from <i>a</i> to <i>b</i> , spacing them evenly
<code>datamash</code>	Perform numeric operations, apply statistical functions, or change formatting on tabular data
<code>vd</code>	VisiData, an interactive tool to explore and operate on tabular data
<code>jp</code>	Draw simple plots from CSV or JSON data
<code>daff</code>	Compare tabular data and find the differences
<code>gnuplot</code>	Utility to plot 2D and 3D graphs
<code>in2csv</code>	Convert various tabular data formats into CSV. Part of the csvkit Python package
<code>sql2csv</code>	Execute SQL queries on a database and output the result as CSV
<code>csvclean</code>	Correct common syntax error on a CSV file
<code>csvcut</code>	Filter data on a CSV file
<code>csvgrep</code>	Find data in specific columns of a CSV file
<code>csvjoin</code>	Join CSV tables (similarly to SQL JOIN)
<code>csvsort</code>	Sort data on a CSV file
<code>csvstack</code>	Stack up rows from CSV files
<code>csvformat</code>	Convert a CSV file to another format
<code>csvjson</code>	Convert a CSV file to JSON format
<code>csvlook</code>	Render a CSV file in the terminal as a fixed width table (compatible with Markdown)
<code>csvpy</code>	Load a CSV file into a Python CLI shell
<code>csvsql</code>	Generate SQL queries from a CSV file and execute them on a database
<code>csvstat</code>	Print statistics for all columns of a CSV file
<code>mlr</code>	Miller, a tool that operates on key-value pair data for querying and reformatting data files in CSV, JSON, and other formats
<code>q</code>	Execute SQL queries against CSV files
<code>textql</code>	Execute SQL queries against CSV files

<code>magick</code>	ImageMagick, a versatile tool to edit, transform, and convert image files
<code>scrot</code>	Take a screenshot
<code>exiftool</code>	Read, write, modify, and delete Exif metadata in image files
<code>exiv2</code>	Read, write, modify, and delete Exif, IPTC, and XMP metadata in image files
<code>zbarimg</code>	Scan an image for barcodes and print any decoded data found
<code>beep</code>	Produce a beep from the machine's speakers
<code>speaker-test</code>	Speaker test tone generator for the ALSA (Advanced Linux Sound Architecture) framework
<code>arecord</code>	Sound recorder for the ALSA soundcard driver
<code>aplay</code>	Sound player for the ALSA soundcard driver
<code>sox</code>	Sound eXchange, "the Swiss Army knife" to read and write audio files
<code>ncmpc</code>	<code>mpd</code> (Music Player Daemon) client with ncurses UI
<code>ncmpcpp</code>	<code>mpd</code> client with improved features with respect to <code>ncmpc</code>
<code>lsdvd</code>	List the contents of a DVD
<code>youtube-dl</code>	Download a video from YouTube

<code>cal</code>	Calendar
<code>banner</code>	Print a text in large letters made of the character #
<code>figlet</code>	Print a text in large letters, in a specific font
<code>toilet</code>	Print a text in large colorful letters, in a specific font
<code>lolcat</code>	Print a text in rainbow coloring
<code>boxes</code>	Draw an ASCII art box around a text
<code>jp2a</code>	Convert a JPG image into ASCII art
<code>cowsay</code>	Display a cow in ASCII art with a message balloon
<code>sl</code>	Useless command that displays a steam locomotive in ASCII art crossing the screen
<code>tesseract</code>	OCR tool to extract text from an image
<code>aspell</code>	Spell checker
<code>dict</code>	Query dictionaries on remote machines via the DICT dictionary protocol
<code>fortune</code>	Print a random aphorism, like those found in fortune cookies
<code>cloc</code>	Count lines of source code
<code>nnn</code>	Terminal file manager
<code>ndn</code>	Necromancer's Dos Navigator, a clone of Norton Commander. With ncurses UI
<code>ipcalc</code>	IP addresses calculator
<code>grepcidr</code>	IP addresses filter against CIDR specifications
<code>on_ac_power</code>	Return 0 (true) if machine is connected to AC power, 1 (false) if on battery. Useful for laptops
<code>pwgen</code>	Random password generator
<code>pwqgen</code>	Random password generator with controllable quality
<code>uuidgen</code>	UUID generator (random or time-based)
<code>haveged</code>	Random number generator using the HAVEGE (Hardware Volatile Entropy Gathering and Expansion) algorithm. Can be run as a daemon to automatically replenish <code>/dev/random</code> whenever the supply of random bits in the random device gets too low
<code>goaccess</code>	Real-time webserver log analyzer with ncurses UI. Can also produce output in HTML format
<code>gotty <i>command</i></code>	Launch a CLI <i>command</i> and show the results in a web page. This is done by running a local web server on port 8080
<code>google-authenticator</code>	Generate a TOTP secret key file for TOTP authentication
<code>latex</code>	LaTeX, the most famous document typesetting system, based on TeX
<code>typst</code>	Markup-based typesetting system, similar to LaTeX

<code>gnome-terminal</code>	GNOME shell terminal GUI
<code>cool-retro-term</code>	Terminal emulator GUI that mimics old cathodic tube screens
<code>kitty</code>	GPU-accelerated terminal emulator GUI with support for images
<code>alacritty</code>	GPU-accelerated terminal emulator GUI
<code>putty</code>	SSH, Telnet, and Rlogin GUI client
<code>gnome-tweaks</code>	GNOME Tweak Tool GUI
<code>fsv</code>	File System Visualizer, a 3D file manager GUI. Open source clone of SGI's <code>f_{sn}</code> for IRIX
<code>notify-send</code>	Program to send desktop notifications
<code>dzen</code>	Messaging and notification tool for X Window
<code>ghostwriter</code>	Text editor for Markdown
<code>qalculate-gtk</code> <code>qalc</code>	Qalculate!, a powerful scientific calculator
<code>wcalc</code>	Tool for the analysis and synthesis of transmission line structures and related components
<code>google-chrome</code> <code>chromium-browser</code>	Google Chrome browser. Can also be run as a headless browser, for machines without a GUI Chromium browser

Red Hat Linux	1995 - 2004	One of the first Linux distros to support ELF binaries.
Red Hat Enterprise Linux (RHEL)	2000 - present	Most used, and de facto standard, commercial Linux distro for production servers in corporate environment. Initially based on Red Hat Linux.
Fedora	2003 - present	Upstream source for RHEL and CentOS Stream. Focuses on innovation and on the early integration of new cutting-edge technologies.
CentOS	2004 - 2021	Free and community-supported Linux distro compiled from RHEL's source code. Downstream of RHEL until 2020, when Red Hat shifted development to CentOS Stream as upstream source for RHEL.
CentOS Stream	2020 - present	Rolling released distro positioned midstream between Fedora and RHEL.
Rocky Linux	2021 - present	Successor to CentOS, created by the original founder of CentOS when Red Hat stopped its development.
AlmaLinux	2021 - present	Free and community-supported Linux distro, created by CloudLinux to replace CentOS when Red Hat stopped its development.
CloudLinux OS	2010 - present	Commercial Linux distro marketed to shared hosting providers and developed by CloudLinux. Based on CentOS.
Oracle Linux (OL)	2006 - present	Free distro compiled from RHEL's source code. Originally called Oracle Enterprise Linux (OEL) .
Scientific Linux (SL)	2004 - 2020	Aimed at scientific environments (labs and universities) and developed originally by Fermilab, CERN, DESY, and ETH Zurich. Derived from RHEL.
Fermi Linux	1998 - 2015	Linux distro created and used at Fermilab. Initially based on RHEL, then on Scientific Linux. Renamed Scientific Linux Fermi in 2004.
Yellow Dog Linux (YDL)	1999 - 2012	Free Linux distro for high performance computing on multi-core machines, first released for Apple Macintosh PowerPC computers. Based on RHEL.
Caldera Network Desktop	1995 - 2002	Early Linux distro based on Red Hat Linux. In 1997 it became Caldera OpenLinux (COL) .
MCC Interim Linux	1992 - 1996	First Linux distro for the general public, released by the University of Manchester.
Yggdrasil Linux/GNU/X (LGX)	1992 - 1995	The first Live CD Linux distro (i.e. usable without installation on the hard disk).
Softlanding Linux System (SLS)	1992 - 1993	First Linux distro to include the X Window System and an extended set of software packages.
Slackware	1993 - 2016	Created as a cleanup of SLS, with focus on design simplicity.
SUSE Linux Enterprise (SLE)	1994 - present	Based on Slackware, and similar to Red Hat Linux. Known as SUSE Linux until 2003.
openSUSE	2005 - present	Fork of SLE aimed at promoting free and open source software.
CRUX	2002 - present	Lightweight Linux distro aimed at experienced users. It uses a BSD-like package management system.
Arch Linux	2002 - present	Focused on design simplicity and minimalism. Inspired by CRUX.
Manjaro	2011 - present	Based on Arch Linux.
Garuda Linux	2020 - present	Based on Arch Linux.
EndeavourOS	2019 - present	Based on Arch Linux. Successor to Antergos .
Gentoo	2000 - present	Distro in which all programs' source code is compiled locally and is customized and optimized for the specific type of computer, resulting in improved performances. Formerly known as Enoch Linux .
Red Flag Linux	1999 - 2020	Linux distro developed in China.
Red Star OS	2008 - present	Official state OS of North Korea, bundled with government spyware. Its UI resembles Microsoft Windows XP (v1 and v2) or Apple macOS (v3 and v4).

Debian	1993 - present	Composed of free and open source software. One of the first Linux distros.
Ubuntu	2004 - present	The most known user-friendly distro, based on Debian. Developed by Canonical. It spawned a number of derivative distros e.g. Lubuntu (lightweight distro with LXQt instead of GNOME), Kubuntu (with KDE), and Xubuntu (with Xfce).
Linux Mint	2006 - present	Based on Ubuntu, offers full multimedia support (codecs, etc).
Pop!_OS	2017 - present	Based on Ubuntu, offers full support for AMD and Nvidia GPUs. Built by computer manufacturer System76 and preinstalled on their systems.
elementary OS	2011 - present	Focused on immediate usability, with a UI resembling Apple macOS. Based on Ubuntu.
Zorin OS	2009 - present	Distro providing a UI that can be customized to resemble Microsoft Windows or Apple macOS. Based on Ubuntu.
Drauger OS	2018 - present	Gaming-focused Linux distribution based on Ubuntu.
Puppy Linux	2003 - 2020	Lightweight and user-friendly distro with minimal memory footprint.
Knoppix	2000 - present	Live CD distro based on Debian.
Kurumin	2003 - 2008	Live CD distro based on Debian, and later on Ubuntu.
Poseidon Linux	2005 - 2011	Based on Kurumin, and later on Ubuntu. Developed by the Rio Grande Federal University, Brazil and the MARUM institute of the University of Bremen, Germany.
Linux Mandrake	1998 - 2011	The first user-friendly Linux distro. Later merged with Conectiva Linux to become Mandriva Linux .
DemoLinux	1998 - 2001	One of the first Live CD Linux distros. Based initially on Linux Mandrake, and later on Debian.
Devuan	2016 - present	Fork of Debian that uses init-like systems instead of systemd.
Damn Small Linux (DSL)	2005 - 2008	Designed to run on older hardware with minimal amounts of RAM. Distributed as a Live CD of about 50 MB in size. Based on Knoppix.
Tiny Core Linux (TCL)	2009 - present	Minimalist Linux distro, about 10 MB in size, based on BusyBox. Created by the developer of Damn Small Linux.
Bayanihan Linux	2003 - 2011	Linux desktop distro developed by the Philippine government. Based originally on Red Hat Linux and Fedora, and later on Debian.
Pardus	2005 - present	Linux desktop distro developed by the Turkish government. Based on Debian.
Astra Linux	2011 - present	Linux distro developed and certified for use within Russian armed forces and intelligence agencies. Based on Debian.
Deepin	2004 - present	Linux distro used mostly in China, criticized for possible breaches of user privacy. Formerly known as Hiweed Linux . Based on Debian.

Kali Linux	2013 - present	The de facto "hacker distro", designed for digital forensics and pentesting. Based on Debian. Rebuild of BackTrack , which was based on Knoppix.
Tsurugi Linux	2018 - present	Designed for DFIR, OSINT, and malware analysis. Based on Ubuntu.
Lightweight Portable Security (LPS)	2007 - 2021	Live CD Linux distro developed by the US Department of Defense and designed to serve as a secure network end node. Renamed Trusted End Node Security (TENS) in 2011.
Tails (The Amnesic Incognito Live System)	2009 - present	Distro focused on privacy and anonymity; runs from a Live USB/DVD, leaves no digital footprint on the machine, and connects to the Internet exclusively via Tor. Based on Debian. Successor to Incognito , which was based on Gentoo.
Kicksecure	2018 - present	Security-hardened distro providing defense-in-depth protection against malicious code. Based on a security reconfiguration of Debian run inside multiple VMs on top of the host OS. Formerly called Hardened Debian .
Whonix	2012 - present	Distro focused on privacy and anonymity, formerly called TorBOX . Consists of two Debian VMs running in parallel inside an isolated network: a "Workstation" and a "Tor Gateway", with network traffic from the Workstation tunneled through the Tor Gateway. Based on Kicksecure.
Qubes OS	2012 - present	Security-focused single-user OS which implements Security by Isolation, running each application in a securely-isolated compartment called qube. A different VM is run (via Xen) for each different domain of trust.

This is a partial list of Linux distributions; more than one thousand Linux distros, either living or defunct, exist.

Locale environment variables	
LANG LANGUAGE	Language, stored in <code>/etc/default/locale</code> . When scripting, it is recommended to set <code>LANG=C</code> because this specifies the minimal locale environment for C translation, and guarantees a standard collation and formats for the execution of scripts
LC_CTYPE	Character classification and case conversion
LC_NUMERIC	Non-monetary numeric formats
LC_TIME	Date and time formats
LC_COLLATE	Alphabetical order
LC_MONETARY	Monetary formats
LC_MESSAGES	Language and encoding of system messages and user input
LC_PAPER	Paper size
LC_NAME	Personal name formats
LC_ADDRESS	Geographic address formats
LC_TELEPHONE	Telephone number formats
LC_MEASUREMENT	Measurement units (metric or others)
LC_IDENTIFICATION	Metadata about locale
LC_ALL	Special variable overriding all others
The values of these locale environment variables are in the format <code>language_territory.encoding</code> e.g. <code>en_US.UTF-8</code> . The list of supported locales is stored in <code>/usr/share/i18n/SUPPORTED</code> .	

<code>export LC_ALL=C</code>	Set the default English locale
<code>locale</code>	Show locale environment variables
<code>locale-gen it_IT.UTF-8</code>	Generate a locale (in this case IT) by compiling a list of locale definition files
<code>apt-get install manpages-it language-pack-it</code> (Debian)	Install a different locale (in this case IT); this affects system messages and manpages
<code>iconv -f ISO-8859-10 filein -t UTF-8 > fileout</code>	Convert a text file from a character set to another
<code>recode cp1251..utf8 file</code>	Convert a text file from a character set to another

ISO/IEC-8859 is a standard for 8-bit encoding of printable characters. The first 256 characters in ISO/IEC-8859-1 (Latin-1) are identical to those in Unicode.

UTF-8 encoding can represent every character in the Unicode set, and is the de facto standard for text containing characters with diacritics, which do not fit in the ASCII 7-bit set. It was designed for backward compatibility with ASCII. UTF-8 encodes a Unicode character into 8, 16, 24, or 32 bits, whatever necessary; a UTF-8 file containing only ASCII characters is identical to an ASCII file.

<code>date</code>	Show current date and time
<code>date -d "9999 days ago"</code>	Calculate a date and show it
<code>date -d "1970/01/01 + 4242"</code>	
<code>date +%F %H:%M:%S</code>	Show current date in the format specified
<code>date +%s</code>	Show current date in Unix time format (i.e. the number of seconds elapsed since 00:00:00 1/1/1970)
<code>date -s "20210104 23:30:00"</code>	Set the date
<code>date 010423302021</code>	Set the date, in the format <i>MMDDhhmmYYYY</i>
<code>timedatectl</code>	Show current date and time
<code>timedatectl set-time 2021-01-04</code>	Set the date
<code>timedatectl set-time 23:30</code>	
<code>timedatectl list-timezones</code>	List all timezones
<code>zdump GMT</code>	Show current date and time in the GMT timezone
<code>tzselect</code>	Set the timezone
<code>tzconfig</code>	
<code>dpkg-reconfigure tzdata</code>	(Debian)
<code>timedatectl set-timezone tz</code>	(Red Hat)
<code>/etc/timezone</code>	(Debian)
<code>/etc/localtime</code>	(Red Hat)
	Timezone
	Timezone. This is a symlink to the appropriate timezone file in <code>/usr/share/zoneinfo/</code>
<code>hwclock --show</code>	Show the hardware clock
<code>hwclock -r</code>	
<code>hwclock --hctosys</code>	Set the system time from the hardware clock
<code>hwclock -s</code>	
<code>hwclock --systohc</code>	Set the hardware clock from system time
<code>hwclock -w</code>	
<code>hwclock --utc</code>	Indicate that the hardware clock is kept in Coordinated Universal Time
<code>hwclock --localtime</code>	Indicate that the hardware clock is kept in local time

NTP (Network Time Protocol) is a protocol used to synchronize machines' clocks in a network, and to keep a machine's clock in sync with Internet timeservers. It uses UDP port 123.

<code>ntpd</code>	NTP daemon
<code>ntpd -q</code>	Synchronize the time once and quit
<code>ntpd -g</code>	Force NTP to start even if clock is off by more than the panic threshold (1000 secs)
<code>ntpd -nqg</code>	Start NTP as a non-daemon, force synchronization of the clock, and quit. The NTP daemon must not be running when this command is launched
<code>ntpdctl</code>	Query and modify the state of the NTP daemon
<code>ntpq -p timeserver</code>	Print the list of peers for the timeserver
<code>ntptrace</code>	Trace from where a NTP server gets its time, following the chain of NTP servers back to the primary time source
<code>ntpdate timeserver</code>	Synchronizes the clock with the specified timeserver
<code>ntpdate -b timeserver</code>	Brutally set the clock, without waiting for it to adjust slowly
<code>ntpdate -q timeserver</code>	Query the timeserver without setting the clock
The <code>ntpdate</code> command is deprecated; to synchronize the clock, <code>ntpd</code> or <code>chronyd</code> should be used instead.	
<code>chronyd</code>	Chrony daemon, a versatile NTP client/server
<code>chronyc</code>	Interactive command line interface for <code>chronyd</code>
<code>chronyc tracking</code>	Display information about the synchronization status of the local machine's clock
<code>chronyc sources</code>	Display information about the time sources
<code>chronyc sourcestats</code>	Display information about drift rate and offset estimation process for each time source

PTP (Precision Time Protocol) is an improvement over NTP, capable of sub-microsecond sync accuracy.


<code>ptp4l</code>	Command-line tool to configure the PTP service
<code>pmc</code>	PTP Management Client
<code>phc2sys</code>	Command-line tool to synchronize two clocks. It is typically used to sync the system clock to a PTP hardware clock (PHC)

```
/etc/syslog.conf
# facility.level          action
*.info;mail.none;authpriv.none /var/log/messages
authpriv.*                /var/log/secure
mail.*                     /var/log/maillog
*.alert                    root
*.emerg                    *
local5.*                   @10.7.7.7
local7.*                   /var/log/boot.log
```

Facility Creator of the message	Level Severity of the message	Action Destination of the message
auth or security† authpriv cron daemon kern lpr mail mark (for syslog internal use) news syslog user uucp local0 ... local17 (custom)	emerg or panic† (highest) alert crit err or error† warning or warn† notice info debug (lowest) none (facility disabled)	file message is written into a log file @host message is sent to a logger server host (via UDP port 514) user1, user2, user3 message is sent to the specified users' consoles * message is sent to all logged in users' consoles
† = deprecated		

Facilities and levels are listed in the manpage `man 3 syslog`.

syslogd	Daemon logging events from user processes
rsyslogd (Ubuntu 14)	
klogd	Daemon logging events from kernel processes
/var/log/messages	Global system logfile
/var/log/dmesg	Kernel ring buffer information
/var/log/kern.log	Kernel log
/var/log/boot.log	Information logged during boot
logger -p auth.info "Message"	Send a message to syslog with facility "auth" and priority "info"
logrotate	Rotate logs. It gzips, renames, and deletes old logfiles according to the configuration files <code>/etc/logrotate.conf</code> and <code>/etc/logrotate.d/*</code> . It is usually scheduled as a daily cron job



<code>~/forward</code>	Mail address(es) to which forward the user's mail, or mail commands
<code>/etc/aliases</code> <code>/etc/mail/aliases</code>	Aliases database for users on the local machine. Each line has syntax <code>alias: user</code>
<code>/var/spool/mail/user</code>	Inbox for <code>user</code> on the local machine
<code>/var/log/mail.log</code> (Debian)	Mail logs
<code>/var/log/maillog</code> (Red Hat)	

<code>mail</code> <code>mailx</code>	Mailclient with advanced commands for non-interactive (batch) use
<code>pine</code>	Mailclient. Obsolete
<code>alpine</code>	Mailclient, a replacement for <code>pine</code>

<code>lsmbx</code>	List the number of messages in a mailbox
--------------------	--

<code>swaks</code>	The Swiss Army's Knife SMTP, a flexible and scriptable SMTP test tool
--------------------	---

<code>mailx -s "Subject" -S smtp="mailserver:25" \</code> <code>user@domain.com < messagefile</code>	Send a mail message to <code>user@domain.com</code> via an external SMTP server <code>mailserver</code>
<code>uuencode binaryfile mail user@domain.com</code>	Send a binary file to <code>user@domain.com</code> (obsolete, not recommended because many mailclients will display the received attachment inline)
<code>mutt -a binaryfile -- user@domain.com < /dev/null</code>	Send a binary file to <code>user@domain.com</code> using the Mutt MUA

Mailbox formats		
mbox	Each mail folder is a single file, storing multiple email messages. Advantages: universally supported; fast search inside a mail folder. Disadvantages: issues with file locking; possible mailbox corruption.	<code>\$HOME/Mail/folder</code>
Maildir	Each mail folder is a directory, and contains the subdirectories <code>/cur</code> , <code>/new</code> , and <code>/tmp</code> . Each email message is stored in its own file with a unique filename ID. The process that delivers an email message writes it to a file in the <code>tmp/</code> directory, and then moves it to <code>new/</code> . The moving is commonly done by hard linking the file to <code>new/</code> and then unlinking the file from <code>tmp/</code> , which guarantees that a MUA will not see a partially written message as it never looks in <code>tmp/</code> . When the MUA finds mail messages in <code>new/</code> it moves them to <code>cur/</code> . Advantages: fast location/retrieval/deletion of a specific mail message; no file locking needed; can be used with NFS. Disadvantages: some filesystems may not efficiently handle a large number of small files; searching text inside all mail messages is slower.	<code>\$HOME/Mail/folder/</code>

SMTP commands		
220 smtp.example.com ESMTP Postfix (server)	HELO xyz.linux.org	Initiate the conversation and identify client host to server
HELO xyz.linux.org (client)		
250 Hello xyz.linux.org, glad to meet you	EHLO xyz.linux.org	Like HELO, but tell server to use Extended SMTP
MAIL FROM: alice@linux.org		
250 Ok	MAIL FROM: alice@linux.org	Specify mail sender
RCPT TO bob@foobar.com	RCPT TO: bob@foobar.com	Specify mail recipient
250 Ok		
RCPT TO carol@quux.net	DATA	Specify data to send. Ended with a dot on a single line
250 Ok		
DATA	QUIT	Disconnect
354 End data with <CR><LF>.<CR><LF>	RSET	
From: Alice <alice@linux.org>	HELP	List all available commands
To: Bob <bob@foobar.com>	NOOP	Empty command
Cc: Carol <carol@quux.net>	VERFY alice@linux.org	Verify the existence of an email address (this command should not be implemented, for security reasons)
Date: Wed, 13 August 2014 18:02:43 -0500		
Subject: Test message	EXPN list@linux.org	Tell the actual delivery address of aliases and mailing lists
This is a test message.		
.		
250 OK id=10jReS-0005kT-Jj		
QUIT		
221 Bye		

SMTP response codes		
first digit	1	Command accepted, but not processed until client sends confirmation
	2	Command successfully completed
	3	Command accepted, but not processed until client sends more information
	4	Command failed due to temporary errors
	5	Command failed due to permanent errors
second digit	0	Syntax error or command not implemented
	1	Informative response in reply to a request for information
	2	Connection response in reply to a data transmission
	5	Status response in reply to a mail transfer operation
third digit	Specifies further the response	
211	System status or help reply	
214	Help message	
220	The server is ready	
221	The server is ending the conversation	
250	The requested action was completed	
251	The specified user is not local, but the server will forward the mail message	
354	Reply to the DATA command. After getting this, start sending the message body	
421	The mail server will be shut down, try again later	
450	The mailbox that you are trying to reach is busy, try again later	
451	The requested action was not done because some error occurred in the mail server	
452	The requested action was not done because the mail server ran out of system storage	
500	The last command contained a syntax error or the command line was too long	
501	The parameters or arguments in the last command contained a syntax error	
502	The last command is not implemented in the mail server	
503	The last command was sent out of sequence	
504	One of the parameters in the last command is not implemented by the server	
550	The mailbox that you are trying to reach cannot be found or you do not have access rights	
551	The specified user is not local, so part of message text will contain a forwarding address	
552	The mailbox that you are trying to reach has run out of space, try again later	
553	The mail address that you specified was not syntactically correct	
554	The mail transaction has failed for unknown causes	

Sendmail is an MTA distributed as a monolithic binary file.

Previous versions used to run SUID `root`, which caused many security problems; recent versions run SGID `smmsp`, the group that has write access on the mail queue.

Sendmail uses `smrsh`, a restricted shell, to run some external programs.

Configuration files (must not be edited by hand):

/etc/mail/	submit.cf	Sendmail local mail transfer configuration file
	sendmail.cf	Sendmail MTA configuration file

`m4 /etc/mail/submit.mc > /etc/mail/submit.cf` Generate a `.cf` configuration file from an editable `.mc` text file

Database files (must not be edited by hand):

/etc/mail/	access.db	Access control file to allow or deny access to systems or users
	local-host-names.db	List of domains that must be considered as local accounts
	virtusertable.db	Map for local accounts, used to distribute incoming email
	mailertable.db	Routing table, used to dispatch emails from remote systems
	domaintable.db	Domain table, used for transitions from an old domain to a new one
	genericstable.db	Map for local accounts, used to specify a different sender for outgoing mail
	genericsdomain.db	Local FQDN

`makemap hash /etc/mail/access.db < /etc/mail/access` Generate a `.db` database file from an editable text file

Temporary mailqueue files (where *nnn* is the Message ID):

/var/spool/mqueue/	dfnnn	Mail body
	qfnnn	Message envelope with headers and routing information
	Qfnnn	Message envelope if abandoned
	hfnnn	Message envelope if held / quarantined by a milter (mail filter)
	tfnnn	Temporary file
	lfnnn	Lock file
	nfnnn	Backup file
	xfnnn	Transcript of delivery attempts

`newaliases` Update the aliases database. Must be run after any change to `/etc/aliases`

`sendmail -bi`

Examine the mail queue

`mailq`

`sendmail -bp`

Run Sendmail in test mode

`sendmail -bt`

`sendmail -q`

Force a queue run

`hoststat`

Print statistics about remote hosts usage

`purgestat`

Clear statistics about remote host usage

`mailstats`

Print statistics about the mailserver

`praliases`

Display email aliases

Exim is a free MTA, distributed under open source GPL license.

<code>/etc/exim.conf</code> <code>/usr/local/etc/exim/configure</code> (FreeBSD)	Exim4 configuration file
<code>exim4 -bp</code>	Examine the mail queue
<code>exim4 -M <i>messageID</i></code>	Attempt delivery of message
<code>exim4 -Mrm <i>messageID</i></code>	Remove a message from the mail queue
<code>exim4 -Mvh <i>messageID</i></code>	See the headers of a message in the mail queue
<code>exim4 -Mvb <i>messageID</i></code>	See the body of a message in the mail queue
<code>exim4 -Mvc <i>messageID</i></code>	See a message in the mail queue
<code>exim4 -qf <i>domain</i></code>	Force a queue run of all queued messages for a <i>domain</i>
<code>exim4 -Rff <i>domain</i></code>	Attempt delivery of all queued messages for a <i>domain</i>
<code>exim4 -bV</code>	Show version and other info
<code>exinext</code>	Give the times of the next queue run
<code>exigrep</code>	Search through Exim logfiles
<code>exicyclog</code>	Rotate Exim logfiles

Postfix is a fast, secure, easy to configure, open source MTA intended as a replacement for Sendmail. It is implemented as a set of small helper daemons, most of which run in a chroot jail with low privileges. The main ones are:

master	Postfix master daemon, always running; starts the other daemons when necessary
nqmgr	Queue manager for incoming and outgoing mail, always running
smtpd	SMTP daemon for incoming mail
smtp	SMTP daemon for outgoing mail
bounce	Manager of bounce messages
cleanup	Daemon that verifies the syntax of outgoing messages before they are handed to the queue manager
local	Daemon that handles local mail delivery
virtual	Daemon that handles mail delivery to virtual users

/var/spool/postfix/	incoming	Incoming queue. All new mail entering the Postfix queue is written here by the cleanup daemon. Under normal conditions this queue is nearly empty.
	active	Active queue. Contains messages ready to be sent. The queue manager places messages here from the incoming queue as soon as they are available.
	deferred	Deferred queue. A message is placed here when all its deliverable recipients are delivered, and delivery failed for some recipients for a transient reason. The queue manager scans this queue periodically and puts some messages back into the active queue to retry sending.
	bounce	Message delivery status report about why mail is bounced (non-delivered mail).
	defer	Message delivery status report about why mail is delayed (non-delivered mail).
	trace	Message delivery status report (delivered mail).

postfix reload	Reload configuration
postconf -e 'mydomain = example.org'	Edit a setting in the Postfix configuration
postconf -l	List supported mailbox lock methods
postconf -m	List supported database types
postconf -v	Increase logfile verbosity
postmap dbtype:textfile	Manage Postfix lookup tables, creating a hashed map file of database type <i>dbtype</i> from <i>textfile</i>
postmap hash:/etc/postfix/transport	Regenerate the transport database
postalias	Convert <i>/etc/aliases</i> into the aliases database file <i>/etc/aliases.db</i>
postsuper	Operate on the mail queue
postqueue	Unprivileged mail queue manager
postqueue -p	Print the queue content
postqueue -f	Flush the queue, attempting delivery of all messages in the queue

<code>/etc/postfix/main.cf</code> Postfix main configuration file	
<code>mydomain = example.org</code>	This system's domain
<code>myorigin = \$mydomain</code>	Domain from which all sent mail will appear to originate
<code>myhostname = foobar.\$mydomain</code>	This system's hostname
<code>inet_interfaces = all</code>	Network interface addresses that this system receives mail on. Value can also be <code>localhost</code> , <code>all</code> , or <code>loopback-only</code>
<code>proxy_interfaces = 1.2.3.4</code>	Network interface addresses that this system receives mail on by means of a proxy or NAT unit
<code>mynetworks = 10.3.3.0/24 !10.3.3.66</code>	Networks the SMTP clients are allowed to connect from
<code>mydestination = \$myhostname, localhost, \$mydomain, example.com, hash:/etc/postfix/otherdomains</code>	Domains for which Postfix will accept received mail. Value can also be a lookup database file e.g. a hashed map
<code>relayhost = 10.6.6.6</code>	Relay host to which Postfix should send all mail for delivery, instead of consulting DNS MX records
<code>relay_domains = \$mydestination</code>	Sources and destinations for which mail will be relayed. Can be empty if Postfix is not intended to be a mail relay
<code>virtual_alias_domains = virtualex.org</code> <code>virtual_alias_maps = /etc/postfix/virtual</code> or <code>virtual_alias_domains = hash:/etc/postfix/virtual</code>	Set up Postfix to handle mail for virtual domains too. The <code>/etc/postfix/virtual</code> file is a hashed map, each line of the file containing the virtual domain email address and the destination real domain email address: <pre> jdoe@virtualex.org john.doe@example.org ksmith@virtualex.org kim.smith @virtualex.org root </pre> The <code>@virtualex.org</code> in the last line is a catch-all specifying that all other email messages to the virtual domain are delivered to the root user on the real domain
<code>mailbox_command = /usr/bin/procmail</code>	Use Procmail as MDA
A line beginning with whitespace or tab is a continuation of the previous line. A line beginning with a # is a comment. A # not placed at the beginning of a line is not a comment delimiter.	

<code>/etc/postfix/master.cf</code> Postfix master daemon configuration file	
<pre> # service type private unpriv chroot wakeup maxproc command + args smtp inet n - - - - smtpd pickup fifo n - - 60 1 pickup cleanup unix n - - - 0 cleanup qmgr fifo n - - 300 1 qmgr rewrite unix - - - - - trivial-rewrite bounce unix - - - - 0 bounce defer unix - - - - 0 bounce flush unix n - - 1000? 0 flush smtp unix - - - - - smtp showq unix n - - - - showq error unix - - - - - error local unix - n n - - local virtual unix - n n - - virtual lmtp unix - - n - - lmtp </pre>	
service	Name of the service
type	Transport mechanism used by the service
private	Whether the service is accessible only by Postfix daemons and not by the whole system. Default is yes
unprivileged	Whether the service is unprivileged i.e. not running as root. Default is yes
chroot	Whether the service is chrooted. Default is yes
wakeup	How often the service needs to be woken up by the master daemon. Default is never
maxproc	Max number of simultaneous processes providing the service. Default is 50
command	Command used to start the service
The - indicates that an option is set to its default value.	

Procmail is a regex-based MDA whose main purpose is to preprocess and sort incoming email messages. It is able to work both with the standard mbox format and the Maildir format.

To have all email processed by Procmail, create a `~/.forward` file with the following content:

```
"|exec /usr/local/bin/procmail || exit 75"
```

`/etc/procmailrc` System-wide recipes

`~/.procmailrc` User's recipes

`procmail -h` List all Procmail flags for recipes

`formail` Utility for email filtering and editing

`lockfile` Utility for mailbox file locking

`mailstat` Utility for generation of reports from Procmail logs

<code>/etc/procmailrc</code> and <code>~/.procmailrc</code> Procmail recipes	
<code>PATH=\$HOME/bin:/usr/bin:/bin:/usr/sbin:/sbin</code> <code>MAILDIR=\$HOME/Mail</code> <code>DEFAULT=\$MAILDIR/Inbox</code> <code>LOGFILE=\$HOME/.procmaillog</code>	Common parameters, nonspecific to Procmail
<code>:0h: or :0:</code> <code>* ^From: .* (alice bob)@foobar\.org</code> <code>\$DEFAULT</code>	Flag: match headers (default) and use file locking (highly recommended when writing to a file or a mailbox in mbox format) Condition: match the header specifying the sender address Destination: default mailfolder
<code>:0:</code> <code>* ^From: .*owner@listserv\.com</code> <code>* ^Subject:.*Linux</code> <code>\$MAILDIR/Geekstuff1</code>	Conditions: match sender address and subject headers Destination: specified mailfolder, in mbox format
<code>:0</code> <code>* ^From: .*owner@listserv\.com</code> <code>* ^Subject:.*Linux</code> <code>\$MAILDIR/Geekstuff2/</code>	Flag: file locking not necessary because using Maildir format Conditions: match sender address and subject headers Destination: specified mailfolder, in Maildir format
<code># Blacklisted by SpamAssassin</code> <code>:0</code> <code>* ^X-Spam-Status: Yes</code> <code>/dev/null</code>	Flag: file locking not necessary because blackholing to <code>/dev/null</code> Condition: match SpamAssassin's specific header Destination: delete the message
<code>:0B:</code> <code>* hacking</code> <code>\$MAILDIR/Geekstuff</code>	Flag: match body of message instead of headers
<code>:0HB:</code> <code>* hacking</code> <code>\$MAILDIR/Geekstuff</code>	Flag: match either headers or body of message
<code>:0:</code> <code>* > 256000</code> <code> /root/myprogram</code>	Condition: match messages larger than 256 Kb Destination: pipe message through the specified program
<code>:0fw</code> <code>* ^From: .*@foobar\.org</code> <code> /root/myprogram</code>	Flags: use the pipe as a filter (modifying the message), and have Procmail wait that the filter finished processing the message
<code>:0c</code> <code>* ^Subject:.*administration</code> <code>! secretary@domain.com</code> <code>:0:</code> <code>\$MAILDIR/Forwarded</code>	Flag: copy the message and proceed with next recipe Destination: forward to specified email address, and (this is ordered by the next recipe) save in the specified mailfolder

Courier is an MTA that provides modules for ESMTP, IMAP, POP3, webmail, and mailing list services in a single framework. To use Courier, it is necessary first to launch the `courier-authlib` service, then launch the desired mail service e.g. `courier-imap` for the IMAP service.

/usr/lib/courier-imap/etc/ or /etc/courier/	imapd	Courier IMAP daemon configuration
	imapd-ssl	Courier IMAPS daemon configuration
	pop3d	Courier POP3 daemon configuration
	pop3d-ssl	Courier POP3S daemon configuration

/usr/lib/courier-imap/share/

Directory for public and private keys

mkimapdcert

Generate a certificate for the IMAPS service

mkpop3dcert

Generate a certificate for the POP3 service

makealiases

Create system aliases in `/usr/lib/courier/etc/aliases.dat`, which is made by processing a `/usr/lib/courier/etc/aliases/system` text file:

```
root      : postmaster
mailer-daemon : postmaster
MAILER-DAEMON : postmaster
uucp      : postmaster
postmaster : admin
```

/usr/lib/courier-imap/etc/pop3d Courier POP configuration file	
ADDRESS=0	Address on which to listen. 0 means all addresses
PORT=127.0.0.1.900,192.168.0.1.900	Port number on which connections are accepted. In this case, accept connections on port 900 on IP addresses 127.0.0.1 and 192.168.0.1
POP3AUTH="LOGIN CRAM-MD5 CRAM-SHA1"	POP authentication advertising SASL (Simple Authentication and Security Layer) capability, with CRAM-MD5 and CRAM-SHA1
POP3AUTH_TLS="LOGIN PLAIN"	Also advertise SASL PLAIN if SSL is enabled
MAXDAEMONS=40	Maximum number of POP3 servers started
MAXPERIP=4	Maximum number of connections to accept from the same IP address
PIDFILE=/var/run/courier/pop3d.pid	PID file
TCPDOPTS="-nodnslookup -noidentlookup"	Miscellaneous <code>couriertcpd</code> options. Should not be changed
LOGGEROPTS="-name=pop3d"	Options for <code>courierlogger</code>
POP3_PROXY=0	Enable or disable proxying
PROXY_HOSTNAME=myproxy	Override value from <code>gethostname()</code> when checking if a proxy connection is required
DEFDOMAIN="@example.com"	Optional default domain. If the username does not contain the first character of <code>DEFDOMAIN</code> , then it is appended to the username. If <code>DEFDOMAIN</code> and <code>DOMAINSEP</code> are both set, then <code>DEFDOMAIN</code> is appended only if the username does not contain any character from <code>DOMAINSEP</code>
POP3DSTART=YES	Flag intended to be read by the system startup script
MAILDIRPATH=Maildir	Maildir directory

/usr/lib/courier-imap/etc/imapd Courier IMAP configuration file	
ADDRESS=0	Address on which to listen. 0 means all addresses
PORT=127.0.0.1.900,192.168.0.1.900	Port number on which connections are accepted. In this case, accept connections on port 900 on IP addresses 127.0.0.1 and 192.168.0.1
AUTHSERVICE143=imap	Authenticate using a different <code>service</code> parameter depending on the connection's port. This only works with authentication modules that use the <code>service</code> parameter, such as PAM
MAXDAEMONS=40	Maximum number of IMAP servers started
MAXPERIP=20	Maximum number of connections to accept from the same IP address
PIDFILE=/var/run/courier/imapd.pid	PID file for <code>couriertcpd</code>
TCPDOPTS="-nodnslookup -noidentlookup"	Miscellaneous <code>couriertcpd</code> options. Should not be changed
LOGGEROPTS="-name=imapd"	Options for <code>courierlogger</code>
DEFDOMAIN="@example.com"	Optional default domain. If the username does not contain the first character of <code>DEFDOMAIN</code> , then it is appended to the username. If <code>DEFDOMAIN</code> and <code>DOMAINSEP</code> are both set, then <code>DEFDOMAIN</code> is appended only if the username does not contain any character from <code>DOMAINSEP</code>
IMAP_CAPABILITY="IMAP4rev1 UIDPLUS \ CHILDREN NAMESPACE THREAD=ORDEREDSUBJECT \ THREAD=REFERENCES SORT QUOTA IDLE"	Specifies what most of the response should be to the <code>CAPABILITY</code> command
IMAP_KEYWORDS=1	Enable or disable custom IMAP keywords. Possible values are: 0 disable keywords 1 enable keywords 2 enable keywords with a slower algorithm
IMAP_ACL=1	Enable or disable IMAP ACL extension
SMTP_CAPABILITY=SMTP	Enable the experimental Simple Mail Access Protocol extensions
IMAP_PROXY=0	Enable or disable proxying
IMAP_PROXY_FOREIGN=0	Proxying to non-Courier servers. Resends the <code>CAPABILITY</code> command after logging in to remote server. May not work with all IMAP clients
IMAP_IDLE_TIMEOUT=60	How often, in seconds, the server should poll for changes to the folder while in IDLE mode
IMAP_CHECK_ALL_FOLDERS=0	Enable or disable server check for mail in every folder
IMAP_UMASK=022	Set the umask of the server process. This value is passed to the <code>umask</code> command. Mostly useful for shared folders, where file permissions of the messages may be important
IMAP_ULIMITD=131072	Set the upper limit of the size of the data segment of the server process, in Kb. This value is passed to the <code>ulimit -d</code> command. Used as an additional safety check to stop potential DoS attacks that exploit memory leaks to exhaust all the available RAM on the server
IMAP_USELOCKS=1	Enable or disable dot-locking to support concurrent multiple access to the same folder. Strongly recommended when using shared folders
IMAP_SHAREDINDEXFILE=\n/etc/courier/shared/index	Index of all accessible folders. This setting should normally not be changed
IMAP_TRASHFOLDERNAME=Trash	Trash folder
IMAP_EMPTYTRASH=Trash:7,Sent:30	Purge folders i.e. delete all messages from the specified folders after the specified number of days
IMAP_MOVE_EXPUNGE_TO_TRASH=0	Enable or disable moving expunged messages to the trash folder (instead of directly deleting them)
HEADERFROM=X-IMAP-Sender	Save the return address (<code>\$SENDER</code>) in the <code>X-IMAP-Sender</code> mail header. This header is added to the sent message, but not in the copy of the message saved in the folder
MAILDIRPATH=Maildir	Mail directory

Dovecot is an open source, security-hardened, fast, and efficient IMAP and POP3 server. It implements its own high-performance mbox mailbox format. By default, it uses PAM authentication. The script `mkcert.sh` can be used to create self-signed SSL certificates.

<code>/etc/dovecot.conf</code> Dovecot configuration file	
<code>base_dir = /var/run/dovecot/</code>	Base directory where to store runtime data
<code>protocols = imap pop3s</code>	Protocols to serve. If Dovecot should use <code>dovecot-auth</code> , this can be set to <code>none</code>
<code>listen = *, [::]</code>	Network interfaces on which to accept connections. In this case, listen to all IPv4 and IPv6 interfaces
<code>disable_plaintext_auth = yes</code>	If yes, disable LOGIN command and all other plaintext authentications unless SSL/TLS is used (LOGINDISABLED capability)
<code>shutdown_clients = yes</code>	If yes, kill all IMAP and POP3 processes when Dovecot master process shuts down; if no, Dovecot can be upgraded without forcing existing client connections to close
<code>log_path = /dev/stderr</code>	Log error messages to a file (in this case, <code>stderr</code>) instead of <code>syslog</code>
<code>info_log_path = /dev/stderr</code>	Log file to use for informational and debug messages. Default value is the same as <code>log_path</code>
<code>syslog_facility = mail</code>	Syslog facility to use, if logging to <code>syslog</code>
<code>login_dir = /var/run/dovecot/login</code>	Directory where the authentication process places authentication UNIX sockets. The login process needs to be able to connect to these sockets
<code>login_chroot = yes</code>	Chroot login process to the <code>login_dir</code>
<code>login_user = dovecot</code>	User for the login process and for access control in the authentication process. This is not the user that will access mail messages
<code>login_process_size = 64</code>	Maximum login process size, in Mb
<code>login_process_per_connection = yes</code>	If yes, each login is processed in its own process (more secure); if no, each login process processes multiple connections (faster)
<code>login_processes_count = 3</code>	Number of login processes to keep for listening for new connections
<code>login_max_processes_count = 128</code>	Maximum number of login processes to create
<code>login_max_connections = 256</code>	Maximum number of connections allowed per each login process. This setting is used only if <code>login_process_per_connection = no</code> ; once the limit is reached, the process notifies master so that it can create a new login process
<code>login_greeting = Dovecot ready.</code>	Greeting message for clients
<code>login_trusted_networks = \n10.7.7.0/24 10.8.8.0/24</code>	Trusted network ranges (usually IMAP proxy servers). Connections from these IP addresses are allowed to override their IP addresses and ports, for logging and authentication checks. <code>disable_plaintext_auth</code> is also ignored for these networks
<code>mbox_read_locks = fcntl</code> <code>mbox_write_locks = dotlock fcntl</code>	Locking methods to use for locking mailboxes in mbox format. Possible values are: dotlock Create <code>mailbox.lock</code> file; oldest and NSF-safe method dotlock_try Same as <code>dotlock</code>, but skip if failing fcntl Recommended; works with NFS too if <code>lockd</code> is used flock May not exist in all systems; doesn't work with NFS lockf May not exist in all systems; doesn't work with NFS
<code>maildir_stat_dirs = no</code>	Option for mailboxes in Maildir format. If no (default), the LIST command returns all entries in the mail directory beginning with a dot; if yes, returns only entries which are directories
<code>dbox_rotate_size = 2048</code> <code>dbox_rotate_min_size = 16</code>	Maximum and minimum file size, in Kb, of a mailbox in dbox format until it is rotated
<code>!include /etc/dovecot/conf.d/*.conf</code>	Include configuration file
<code>!include_try /etc/dovecot/extra.conf</code>	Include optional configuration file, and do not report an error if file is not found

/etc/dovecot.conf	Dovecot configuration file
<pre>mail_location = \ mbox:~/mail:INBOX=/var/spool/mail/%u or mail_location = maildir:~/Maildir</pre>	Mailbox location, in mbox or Maildir format. Variables: %u username %n user part in <i>user@domain</i>, same as %u if there is no domain %d domain part in <i>user@domain</i>, empty if there is no domain %h home directory
<pre>namespace shared { separator = / prefix = shared/%u/ location = maildir:%h/Maildir:\ INDEX=~/Maildir/shared/%u inbox = no hidden = no subscriptions = no list = children }</pre>	Definition of a shared namespace, for accessing other users' mailboxes that have been shared. Private namespaces are for users' personal emails. Public namespaces are for shared mailboxes managed by root user Hierarchy separator to use. It should be the same for all namespaces, and depends on the underlying mail storage format Prefix required to access this namespace; must be different for each. In this case, mailboxes are visible under <i>shared/user@domain/</i>; the variables %n, %d, and %u are expanded to the destination user Mailbox location for other users' mailboxes; it is in the same format as <i>mail_location</i> which is also the default for it. %variable and ~/ expand to the logged in user's data; %%variable expands to the destination user's data Define whether this namespace contains the INBOX. Note that there can be only one INBOX across all namespaces Define whether the namespace is hidden i.e. not advertised to clients via NAMESPACE extension Namespace handles its own subscriptions; if set to no, the parent namespace handles them and Dovecot uses the default namespace for saving subscriptions. If <i>prefix</i> is empty, this should be set to yes Show the mailboxes under this namespace with LIST command, making the namespace visible for clients that do not support the NAMESPACE extension. In this case, lists child mailboxes but hide the namespace prefix; list the namespace only if there are visible shared mailboxes
<pre>mail_uid = 666 mail_gid = 666</pre>	UID and GID used to access mail messages
<pre>mail_privileged_group = mail</pre>	Group to enable temporarily for privileged operations. Currently this is used only with INBOX when its initial creation or a dotlocking fails
<pre>mail_access_groups = tmpmail</pre>	Supplementary groups to with grant access for mail processes. Used typically to set up access to shared mailboxes
<pre>lock_method = fcntl</pre>	Locking method for index files. Can be <i>fcntl</i>, <i>flock</i>, or <i>dotlock</i>
<pre>first_valid_uid = 500 last_valid_uid = 0</pre>	Valid UID range for users; default is 500 and above. This makes sure that users cannot login as daemons or other system users. Denying root login is hardcoded to Dovecot and cannot be bypassed
<pre>first_valid_gid = 1 last_valid_gid = 0</pre>	Valid GID range for users; default is non-root. Users with invalid primary GID are not allowed to login
<pre>max_mail_processes = 512</pre>	Maximum number of running mail processes. When this limit is reached, new users are not allowed to login
<pre>mail_process_size = 256</pre>	Maximum mail process size, in Mb
<pre>valid_chroot_dirs =</pre>	List of directories under which chrooting is allowed for mail processes
<pre>mail_chroot =</pre>	Default chroot directory for mail processes. Usually not needed as Dovecot does not allow users to access files outside their mail directory
<pre>mailbox_idle_check_interval = 30</pre>	Minimum time, in seconds, to wait between mailbox checks. When the IDLE command is running, mailbox is checked periodically for new mails or other changes

<code>/etc/dovecot.conf</code> Dovecot configuration file	
<pre>protocol pop3 { listen = *:110 login_executable = /usr/libexec/dovecot/pop3-login mail_executable = /usr/libexec/dovecot/pop3 pop3_no_flag_updates = no pop3_lock_session = no pop3_uidl_format = %08Xu%08Xv }</pre>	Block with options for the POP3 protocol Network interfaces on which to accept POP3 connections Location of the POP3 login executable Location of the POP3 mail executable If set to no, do not try to set mail messages non-recent or seen with POP3 sessions, to reduce disk I/O. With Maildir format do not move files from <code>new/</code> to <code>cur/</code>; with mbox format do not write <code>Status-</code> headers Defines whether to keep the mailbox locked for the whole POP3 session POP3 UIDL (Unique Mail Identifier) format to use
<pre>protocol imap { listen = *:143 ssl_listen = *:993 login_executable = /usr/libexec/dovecot/imap-login mail_executable = /usr/libexec/dovecot/imap mail_max_userip_connections = 10 imap_idle_notify_interval = 120 }</pre>	Block with options for the IMAP protocol Network interfaces on which to accept IMAP and IMAPS connections Location of the IMAP login executable Location of the IMAP mail executable Maximum number of IMAP connections allowed for a user from each IP address Waiting time, in seconds, between "OK Still here" notifications when client is IDLE
<code>ssl = yes</code>	SSL/TLS support. Possible values are <code>yes</code> , <code>no</code> , <code>required</code>
<code>ssl_cert_file = /etc/ssl/certs/dovecot-cert.pem</code>	Location of the SSL certificate
<code>ssl_key_file = /etc/ssl/private/dovecot-key.pem</code>	Location of private key
<code>ssl_key_password = p4ssw0rd</code>	Password of private key, if it is password-protected. Since <code>/etc/dovecot.conf</code> is usually world-readable, it is better to place this setting into a root-owned 0600 file instead and include it via the setting <code>!include_try /etc/dovecot/dovecot-passwd.conf</code> . Alternatively, Dovecot can be started with <code>dovecot -p p4ssw0rd</code>
<code>ssl_ca_file = /etc/dovecot/cafile.pem</code>	List of trusted SSL certificate authorities. This file contains CA certificates followed by CRLs
<code>ssl_verify_client_cert = yes</code>	Request client to send a certificate
<code>ssl_cipher_list = ALL:!LOW:!SSLv2</code>	List of SSL ciphers to use
<code>verbose_ssl = yes</code>	Show protocol level SSL errors

<code>/etc/dovecot.conf</code> Dovecot configuration file	
<code>auth_executable = /usr/libexec/dovecot/dovecot-auth</code>	Location of the authentication executable
<code>auth_process_size = 256</code>	Max authentication process size, in Mb
<code>auth_username_chars = abcde...VWXYZ01234567890.-_@</code>	List of allowed characters in the username. If the username entered by the user contains a character not listed in here, the login automatically fails. This is to prevent a user exploiting any potential quote-escaping vulnerabilities with SQL/LDAP databases
<code>auth_realms =</code>	List of realms for SASL authentication mechanisms that need them. If empty, multiple realms are not supported
<code>auth_default_realm = example.org</code>	Default realm/domain to use if none was specified
<code>auth_anonymous_username = anonymous</code>	Username to assign to users logging in with ANONYMOUS SASL mechanism
<code>auth_verbose = no</code>	Defines whether to log unsuccessful authentication attempts and the reasons why they failed
<code>auth_debug = no</code>	Define whether to enable more verbose logging (e.g. SQL queries) for debugging purposes
<code>auth_failure_delay = 2</code>	Delay before replying to failed authentications, in seconds
<pre>auth default { mechanisms = plain login cram-md5 passdb passwd-file { args = /etc/dovecot/deny deny = yes } passdb pam { args = cache_key=%u%r dovecot } passdb passwd { blocking = yes args = } passdb shadow { blocking = yes args = } passdb bsdauth { cache_key = %u args = } passdb sql { args = /etc/dovecot/dovecot-sql.conf } passdb ldap { args = /etc/dovecot/dovecot-ldap.conf } socket listen { master { path = /var/run/dovecot/auth-master mode = 0600 user = group = } client { path = /var/run/dovecot/auth-client mode = 0660 } } }</pre>	Accepted authentication mechanisms Deny login to the users listed in <code>/etc/dovecot/deny</code> (this file contains one user per line) PAM authentication block. Enables authentication matching (username and remote IP address) for PAM System users e.g. NSS or <code>/etc/passwd</code> Shadow passwords for system users, e.g. NSS or <code>/etc/passwd</code> PAM-like authentication for OpenBSD SQL database LDAP database Export the authentication interface to other programs. Master socket provides access to userdb information, and is typically used to give Dovecot's local delivery agent access to userdb so it can find mailbox locations. The default user/group is the one who started <code>dovecot-auth</code> (i.e. root). The client socket is generally safe to export to everyone. Typical use is to export it to the SMTP server so it can do SMTP AUTH lookups using it

FTP (File Transfer Protocol) is a client-server unencrypted protocol for file transfer. It uses TCP port 20 for data transfer and TCP port 21 for control.

Its secured version is **FTPS (FTP Secure aka FTP-SSL)**, an extension to FTP with support for TLS, which uses TCP port 989 for data transfer and TCP port 990 for control.

Another secure alternative for file transfer is **SFTP (Secure File Transfer Protocol aka SSH File Transfer Protocol)**, an extension to SSH which has nothing to do with the original FTP.

TFTP (Trivial File Transfer Protocol) is a simple version of FTP which allows a client to transfer a file from or to a remote host. It lacks any login or access control mechanism, and is almost exclusively used in LANs, where one of its primary uses is in the early stages of nodes booting. It uses UDP port 69.

FTP can operate either in active or in passive mode:

Active mode (default)

1. Client connects to FTP server on port 21 (control channel) and sends second unprivileged port number
2. Server acknowledges
3. Server connects from port 20 (data channel) to client's second unprivileged port number
4. Client acknowledges

Passive mode (more protocol-compliant, because it is the client, not the server, that initiates the second connection)

1. Client connects to FTP server on port 21 and requests passive mode via the PASV command
2. Server acknowledges and sends unprivileged port number via the PORT command
3. Client connects to server's unprivileged port number
4. Server acknowledges

<code>/var/log/xferlog.log</code>	Transfer log file from the FTP server, containing full information about the transferred files
<code>ftp</code>	Standard FTP client
<code>lftp</code>	Sophisticated FTP client with support for HTTP and BitTorrent
<code>tftp</code>	TFTP client
<code>ftpd</code>	Main daemon for the DARPA FTP server
<code>pure-ftpd</code>	Main daemon for Pure-FTP, a free and easy-to-use FTP server
<code>pure-ftpwho</code>	Show clients connected to the Pure-FTP server
<code>pure-mrtginfo</code>	Show connections to the Pure-FTP server as a MRTG graph
<code>pure-statsdecode</code>	Show Pure-FTP log data
<code>pure-pw</code>	Manage Pure-FTP virtual accounts
<code>pure-pwconvert</code>	Convert the system user database to a Pure-FTP virtual accounts database
<code>pure-quotacheck</code>	Manage Pure-FTP quota database
<code>pure-uploadscript</code>	Run a command on the Pure-FTP server to process an uploaded file

Very Secure FTP is a hardened and high-performance FTP implementation. The `vsftpd` daemon operates with multiple processes that run as a non-privileged user in a chrooted jail.

/etc/vsftpd/vsftpd.conf	Very Secure FTP server configuration file
listen=NO	Run <code>vsftpd</code> in standalone mode (i.e. not via <code>inetd</code>)?
local_enable=YES	Allow local system users (i.e. in <code>/etc/passwd</code>) to log in?
chroot_local_user=YES	Chroot local users in their home directory?
write_enable=YES	Allow FTP commands that write on the filesystem (i.e. <code>STOR</code> , <code>DELE</code> , <code>RNFR</code> , <code>RNTO</code> , <code>MKD</code> , <code>RMD</code> , <code>APPE</code> , and <code>SITE</code>)?
anonymous_enable=YES	Allow anonymous logins? If yes, <code>anonymous</code> and <code>ftp</code> are accepted as logins
anon_root=/var/ftp/pub	Directory to go after anonymous login
anon_upload_enable=YES	Allow anonymous uploads?
chown_uploads=YES	Change ownership of anonymously uploaded files?
chown_username=ftp	User to whom set ownership of anonymously uploaded files
anon_world_readable_only=NO	Allow anonymous users to only download world-readable files?
pam_service_name=vsftpd	PAM configuration file (located in the <code>/etc/pam.d/</code> directory)
ssl_enable=YES	Enable SSL?
force_local_data_ssl=NO	Encrypt local data?
force_local_logins_ssl=YES	Force encrypted authentication?
allow_anon_ssl=YES	Allow anonymous users to use SSL?
ssl_tlsv1=YES ssl_tlsv2=NO ssl_tlsv3=NO	Allowed SSL/TLS versions
rsa_cert_file=/etc/pki/tls/certs/vsftpd.pem	Location of certificate file
rsa_private_key_file=/etc/pki/tls/certs/vsftpd.pem	Location of private key file

In Linux, printers are managed by `cupsd`, the **CUPS (Common Unix Printing System)** daemon. Printers are administered via a web interface on the URL `http://localhost:631`.

<code>/etc/cups/cupsd.conf</code>	CUPS configuration file
<code>/etc/cups/printers.conf</code>	Database of available local CUPS printers
<code>/etc/printcap</code>	Database of printer capabilities, for old printing applications
<code>/var/spool/cups/</code>	Printer spooler for data awaiting to be printed
<code>/var/log/cups/error_log</code>	CUPS error log
<code>/var/log/cups/page_log</code>	Information about printed pages
<code>/etc/init.d/cupsys start</code>	Start the CUPS service
<code>gnome-cups-manager</code>	Run the CUPS Manager graphical application
<code>cupsenable printer0</code>	Enable a CUPS printer
<code>cupsdisable printer0</code>	Disable a CUPS printer
<code>cupsaccept printer0</code>	Accept a job sent on a printer queue
<code>cupsreject -r "Message" printer0</code>	Reject a job sent on a printer queue, with an informational message
<code>cupstestppd LEXC510.ppd</code>	Test the conformance of a PPD file to the format specification
<code>cupsaddsmb printer0</code>	Export a printer to Samba (for use with MS Windows clients)
<code>cups-config --cflags</code>	Show the necessary compiler options
<code>cups-config --datadir</code>	Show the default CUPS data directory
<code>cups-config --ldflags</code>	Show the necessary linker options
<code>cups-config --libs</code>	Show the necessary libraries to link to
<code>cups-config --serverbin</code>	Show the default CUPS binaries directory that stores filters and backends
<code>cups-config --serverroot</code>	Show the default CUPS configuration file directory
<code>lpstat</code>	Show CUPS status information
<code>lpadmin</code>	Administer CUPS printers
<code>lpadmin -p printer0 -P LEXC750.ppd</code>	Specify a PPD (Adobe PostScript Printer Description) file to associate to a printer
<code>lp -d printer0 file</code>	Print a file on the specified printer
<code>lpq</code>	View the default print queue
<code>lpq -P printer0</code>	View a specific print queue
<code>lpq user</code>	View the print queue of a specific user
<code>lprm -P printer0 jobnumber</code>	Delete a specific job from a printer queue
<code>lprm -P printer0 user</code>	Delete all jobs from a specific user from a printer queue
<code>lprm -P printer0 -</code>	Delete all jobs from a printer queue
<code>lpc</code>	Manage print queues
<code>a2ps file.txt</code>	Convert a text file to PostScript
<code>ps2pdf file.ps</code>	Convert a file from PostScript to PDF
<code>mpage file.ps</code>	Print a PostScript document on multiple pages per sheet on a PostScript printer
<code>gv file.ps</code>	View a PostScript document. Derivation of GhostView

IPv4 addressing					
		Address range	Prefix	Number of addresses	Reference
Classful	Class A (Unicast)	0.0.0.0 – 127.255.255.255 first octet: 0XXX XXXX	/8	128 networks × 16,777,216 addresses	RFC 791
	Class B (Unicast)	128.0.0.0 – 191.255.255.255 first octet: 10XX XXXX	/16	16,384 networks × 65,536 addresses	RFC 791
	Class C (Unicast)	192.0.0.0 – 223.255.255.255 first octet: 110X XXXX	/24	2,097,152 networks × 256 addresses	RFC 791
	Class D (Multicast)	224.0.0.0 – 239.255.255.255 first octet: 1110 XXXX	/4	268,435,456	RFC 3171
	Class E (Experimental)	240.0.0.0 – 255.255.255.255 first octet: 1111 XXXX	/4	268,435,456	RFC 1166
Private	Private Class A	10.0.0.0 – 10.255.255.255	10.0.0.0/8	16,777,216	RFC 1918
	Private Class B	172.16.0.0 – 172.31.255.255	172.16.0.0/12	1,048,576	RFC 1918
	Private Class C	192.168.0.0 – 192.168.255.255	192.168.0.0/16	65,536	RFC 1918
Reserved	Source	0.0.0.0 – 0.255.255.255	0.0.0.0/8	16,777,216	RFC 1700
	Loopback	127.0.0.0 – 127.255.255.255	127.0.0.0/8	16,777,216	RFC 1700
	Autoconf	169.254.0.0 – 169.254.255.255	169.254.0.0/16	65,536	RFC 3330
	TEST-NET	192.0.2.0 – 192.0.2.255	192.0.2.0/24	256	RFC 3330
	6to4 relay anycast	192.88.99.0 – 192.88.99.255	192.88.99.0/24	256	RFC 3068
	Device benchmarks	198.18.0.0 – 198.19.255.255	198.18.0.0/15	131,072	RFC 2544

An IPv4 address is 32-bit long, and is represented divided in four octets (dotted-quad notation), e.g. 193.22.33.44.

There are approximately 4×10^9 total possible IPv4 addresses.

IPv4 classful addressing is obsolete and has been replaced by CIDR (Classless Inter-Domain Routing).

IPv6 addressing	
Unicast	64-bit network prefix ($\geq$ 48-bit routing prefix + $\leq$ 16-bit subnet id) + 64-bit interface identifier A 48-bit MAC address is transformed into a 64-bit EUI-64 by inserting ff:fe in the middle. A EUI-64 is then transformed into an IPv6 interface identifier by inverting the 7 th most significant bit.
Link-local	fe80:0000:0000:0000 + 64-bit interface identifier
Multicast	ff + 4-bit flag + 4-bit scope field + 112-bit group ID

An IPv6 address is 128-bit long, and is represented divided in eight 16-bit groups (4 hex digits).

Leading zeros in each group can be deleted. A single chunk of one or more adjacent 0000 groups can be deleted.
e.g. 2130:0000:0000:0000:0007:0040:15bc:235f which can also be written as 2130::7:40:15bc:235f.

There are approximately 3×10^{38} total possible IPv6 addresses.

The IANA (Internet Assigned Numbers Authority) manages the allocation of IPv4 and IPv6 addresses, assigning large blocks to RIRs (Regional Internet Registries) which in turn allocate addresses to ISPs (Internet Service Providers) and other local registries. These address blocks can be searched via a WHOIS query to the appropriate RIR, which is:

AFRINIC	for Africa
ARIN	for US, Canada, and Antarctica
APNIC	for Asia and Oceania
LACNIC	for Latin America
RIPE NCC	for Europe, Middle East, and Russia

VLSM chart - Last octet subnetting (CIDR notation)						
Prefix: /24 Netmask: .0 00000000 1 subnet 254 hosts each 254 total hosts	Prefix: /25 Netmask: .128 10000000 2 subnets 126 hosts each 252 total hosts	Prefix: /26 Netmask: .192 11000000 4 subnets 62 hosts each 248 total hosts	Prefix: /27 Netmask: .224 11100000 8 subnets 30 hosts each 240 total hosts	Prefix: /28 Netmask: .240 11110000 16 subnets 14 hosts each 224 total hosts	Prefix: /29 Netmask: .248 11111000 32 subnets 6 hosts each 192 total hosts	Prefix: /30 Netmask: .252 11111100 64 subnets 2 hosts each 128 total hosts
0	.0	.0	.0	.0	.0	.0
						.4
					.8	.8
						.12
				.16	.16	.16
						.20
					.24	.24
						.28
			.32	.32	.32	.32
						.36
					.40	.40
						.44
				.48	.48	.48
						.52
					.56	.56
						.60
		.64	.64	.64	.64	.64
						.68
					.72	.72
						.76
				.80	.80	.80
						.84
					.88	.88
						.92
			.96	.96	.96	.96
						.100
					.104	.104
						.108
				.112	.112	.112
						.116
					.120	.120
						.124
	0.13	.128	.128	.128	.128	.128
						.132
					.136	.136
						.140
				.144	.144	.144
						.148
					.152	.152
						.156
			.160	.160	.160	.160
						.164
					.168	.168
						.172
				.176	.176	.176
						.180
					.184	.184
						.188
		.192	.192	.192	.192	.192
						.196
					.200	.200
						.204
				.208	.208	.208
						.212
					.216	.216
						.220
			.224	.224	.224	.224
						.228
					.232	.232
						.236
				.240	.240	.240
						.244
					.248	.248
						.252


Each block of a column identifies a subnet.

The network address of the subnet is the number shown inside a block.

The broadcast address of the subnet is the network address of the block underneath -1 or, for the bottom block, .255.

The range of valid host addresses of the subnet is [network address +1 — broadcast address -1], inclusive.

ISO/OSI and TCP/IP protocol stack models				
Layer	ISO/OSI	TCP/IP	Standards	Data transmission unit
7	Application	Application	HTTP, SMTP, POP, etc.	Message
6	Presentation			
5	Session			
4	Transport	Transport	TCP, UDP	Segment (TCP), datagram (UDP)
3	Network	Internet	IPv4, IPv6, ICMP, etc.	Packet
2	Data Link	Network Access	Ethernet, Wi-Fi, etc.	Frame
1	Physical			



Network bandwidth is the maximum rate of how much data can be transmitted through the network per amount of time. It is usually measured in Mbps (Megabits per second).

Network latency is the rate of how long a signal takes to be transmitted to its destination and back, i.e. its RTT (Round-Trip Time). It is usually measured in ms (milliseconds). It is also informally called **ping**, as the `ping` command is commonly used to measure the RTT. High latency is informally called **lag**. The difference in latency amongst different packets is called **packet delay variation (PDV)** or, informally, **jitter**.

Most common wireless standards				
IEEE standard	Known as	Frequency (GHz)	Max bandwidth (Mbps)	Max range (m)
802.11	Wi-Fi	2.4	2	100
802.11a		5	54	100
802.11b		2.4	11	150
802.11g		2.4	54	150
802.11n		2.4, 5	54, 600	250
802.15.1	Bluetooth	2.4	50	10 - 250
802.16	WiMax	2 - 11	1000	10000

Wireless transmission techniques	
Direct-Sequence Spread Spectrum (DSSS)	Spread-spectrum modulation technique that modulates the original data with a pseudorandom bit sequence (spreading sequence). It is used to reduce signal interference.
Frequency-Hopping Spread Spectrum (FHSS)	Radio transmission technique consisting in rapidly changing the carrier frequency amongst different frequencies, in sync between transmitter and receiver. It is used to reduce signal interference, avoid eavesdropping, and allow code-division multiple access (CDMA) communications.
Orthogonal Frequency-Division Multiplexing (OFDM)	Digital multi-carrier modulation technique which uses multiple orthogonal subcarrier signal frequencies to transmit data, mapping information on the changes in the carrier phase, frequency, or amplitude. It is used to cope with severe channel conditions.
Multiple-Input Multiple-Output Orthogonal Frequency-Division Multiplexing (MIMO-OFDM)	Access mode for 4G and 5G broadband wireless communications. It is used to increase spectral efficiency and reduce signal interference.

Wireless encryption algorithms	
WEP (Wired Equivalent Privacy) IEEE 802.11	WEP uses a pre-shared key with a length of 40, 104, or 232 bits, with a random 24-bit IV (Initialization Vector) added to the key. A CRC-32 checksum is computed on the data and added to it as ICV (Integrity Check Value). WEP key and IV are fed to the RC4 stream cipher to generate a key stream, which is XORed with the data and ICV to obtain the encrypted data. WEP is insecure because of the short length of the IV, which leads to IV reuse; furthermore, the WEP standard does not even require a different IV for each packet. Key reuse in a stream cipher is bad practice and leads to weak encryption. A weak IV may even allow to deduce the WEP pre-shared key. In case of an IV collision, it is possible to reconstruct the RC4 key stream from the IV and the packet's decrypted payload. Therefore, WEP does not provide cryptographic integrity protection of a packet, and is now obsolete. Some wireless APs use LEAP (Lightweight Extensible Authentication Protocol), a Cisco proprietary version of the EAP authentication method for WLANs. LEAP can use either dynamic WEP keys (keys that change very often to minimize cracking exposure) or TKIP. LEAP uses either the MS-CHAP or the EAP-FAST authentication protocol. However, WEP with LEAP is still considered vulnerable.
WPA (Wi-Fi Protected Access) draft IEEE 802.11i	In WPA, the TKIP (Temporal Key Integrity Protocol) feeds a 128-bit temporal key and a 64-bit MIC (Message Integrity Check) to the RC4 stream cipher to obtain the encrypted data. It uses the CRC-32 checksum algorithm strengthened by the use of Michael MIC codes. IV size is 48 bits. TKIP adds a rekeying mechanism to provide fresh encryption and integrity keys, changing temporal keys every 10000 packets in sync between Access Point and client.
WPA2 (Wi-Fi Protected Access II) IEEE 802.11i	WPA2 is encrypted using CCMP (Counter Mode CBC-MAC Protocol), which utilizes AES encryption. IV size is 48 bits. WPA2-Personal uses a PSK (Pre-Shared Key). The Access Point encrypts the data using a 128-bit key derived from a passphrase with length from 8 to 63 characters. Encryption keys are unique for each client, and change frequently. WPA2-Enterprise uses centralized client authentication via 802.1X, either EAP (Extensible Authentication Protocol) or RADIUS (Remote Authentication Dial-In User Service). A TLS-encapsulated secured version of EAP, called PEAP (Protected Extensible Authentication Protocol), is also available. After PSK or 802.1X authentication, a shared secret key called PMK (Pairwise Master Key) is generated, and is validated through a four-way handshake between wireless client and Access Point: 1. AP sends a nonce to the client, which uses it to build the PTK (Pairwise Transient Key) 2. The client sends a nonce and a MIC to the AP 3. The AP builds and sends the GTK (Group Temporal Key) with another MIC to the client 4. The client acknowledges reception to the AP
Wireless encryption can be used in conjunction with other security measures such as SSID cloaking (security by obscurity) and MAC address filtering (whitelisting), which however are not very effective.	

Most common well-known ports	
Port number	Service
13 TCP	Daytime Protocol
20 TCP	FTP (data)
21 TCP	FTP (control)
22 TCP	SSH
23 TCP	Telnet
25 TCP	SMTP
53 TCP/UDP	DNS
67 UDP	BOOTP/DHCP (server)
68 UDP	BOOTP/DHCP (client)
69 UDP	TFTP
80 TCP	HTTP
88 TCP	Kerberos
110 TCP	POP3
119 TCP	NNTP
123 UDP	NTP
135 TCP/UDP	Microsoft RPC
137 TCP/UDP	Microsoft NetBIOS Name Service / WINS
138 TCP/UDP	Microsoft NetBIOS Datagram Service
139 TCP/UDP	Microsoft NetBIOS Session Service
143 TCP	IMAP
161 UDP	SNMP
162 TCP/UDP	SNMP Trap
389 TCP/UDP	LDAP
443 TCP	HTTPS (HTTP over SSL/TLS)
445 TCP/UDP	Microsoft SMB
465 TCP	SMTP over SSL
500 UDP	IPSec ISAKMP / IKE
514 UDP	Syslog
515 TCP/UDP	Line Printer Daemon
901 TCP	Samba SWAT
993 TCP	IMAPS (IMAP over SSL)
995 TCP	POP3S (POP3 over SSL)
4500 UDP	IPSec NAT Traversal

1-1023: privileged ports, used server-side

1024-65535: unprivileged ports, used client-side

The file `/etc/services` lists all well-known ports.

Many network services are run by `xinetd`, the Extended Internet services daemon, rather than standalone. `xinetd` operates as a super server, listening on all service ports defined in its configuration, and upon receiving a connection request it starts the appropriate service.

ip a ip addr ip addr show ifconfig -a	Display configuration of all network interfaces
ip link show eth0 ifconfig eth0	Display configuration of <code>eth0</code>
ip addr add dev eth0 10.1.1.3/24 ifconfig eth0 10.1.1.3 netmask 255.255.255.0 broadcast 10.1.1.255	Configure IP address of <code>eth0</code>
ifconfig eth0 hw ether 45:67:89:ab:cd:ef	Configure MAC address of <code>eth0</code>
ip link set eth0 up ifconfig eth0 up ifup eth0	Activate <code>eth0</code>
ip link set eth0 down ifconfig eth0 down ifdown eth0	Deactivate <code>eth0</code>
dhclient eth0 pump -i eth0 dhcpcd eth0 (SUSE)	Request an IP address via DHCP
ip neigh arp -a	Show the ARP cache table (containing mappings of MAC to IP addresses)
ip neigh show 10.1.1.4 arp 10.1.1.4	Show the ARP cache entry for a host
ip neigh add 10.1.1.5 lladdr 01:23:45:67:89:ab dev eth0 arp -s 10.1.1.5 01:23:45:67:89:ab	Add a new ARP entry for a host
ip neigh del 10.1.1.5 dev eth0 arp -d 10.1.1.5	Delete an ARP entry
ip neigh flush all	Delete the ARP table for all interfaces
/etc/init.d/networking restart (Debian) /etc/init.d/network restart (Red Hat)	Restart network services
hostname	Display the hostname
hostname -f	Display the FQDN (Fully Qualified Domain Name)
hostname <i>name</i> hostnamectl set-hostname --static " <i>name</i> " (Red Hat)	Set the hostname
hostnamectl (Red Hat)	Display the hostname, OS, and other information
ethtool <i>option device</i> ethtool eth0	Query or control network driver and hardware settings View hardware settings of <code>eth0</code>
nsdomainname	Display the DNS domain name
domainname nisdomainname ypdomainname	Display or set the NIS/YP domain name. The Network Information Service (NIS) , originally called Yellow Pages (YP) , is a distributed directory service used to manage and distribute system configuration data, hostname information, user account information, and other data across a computer network

<code>/etc/hosts</code>	Mappings between IP addresses and hostnames, for name resolution <pre>127.0.0.1 localhost.localdomain localhost 10.2.3.4 myhost.domain.org myhost</pre>
<code>/etc/nsswitch.conf</code>	Sources for various system library lookup functions <pre>passwd: files nisplus nis shadow: files nisplus nis group: files nisplus nis hosts: files dns nisplus nis</pre>
<code>/etc/host.conf</code>	Sources for name resolution, for systems before glibc2. Obsolete, superseded by <code>/etc/nsswitch.conf</code> <pre>order hosts,bind multi on</pre>
<code>/etc/resolv.conf</code>	Domain names that must be appended to bare hostnames, and DNS servers that will be used for name resolution <pre>search domain1.org domain2.org nameserver 192.168.3.3 nameserver 192.168.4.4</pre>
<code>/etc/networks</code>	Mappings between network addresses and names <pre>loopback 127.0.0.0 mylan 10.2.3.0</pre>
<code>/etc/services</code>	List of service TCP/UDP port numbers
<code>/etc/protocols</code>	List of available protocols
<code>/sys/class/net</code>	List of all network interfaces in the system

/etc/sysconfig/network (Red Hat)

Network configuration file

```
ADDRESS=10.2.3.4
NETMASK=255.255.255.0
GATEWAY=10.2.3.254
HOSTNAME=mylinuxbox.example.org
NETWORKING=yes
```

/etc/sysconfig/network-scripts/ifcfg-eth0 (Red Hat)

Configuration file for eth0.

This file is read by the ifup and ifdown scripts

```
DEVICE=eth0
TYPE=Ethernet
HWADDR=AA:BB:CC:DD:EE:FF
BOOTPROTO=none
ONBOOT=yes
NM_CONTROLLED=no
IPADDR=10.2.3.4
NETMASK=255.255.255.0
GATEWAY=10.2.3.254
DNS1=8.8.8.8
DNS2=4.4.4.4
USERCTL=no
```

/etc/sysconfig/network-scripts/ifcfg-eth0:0 (Red Hat)
/etc/sysconfig/network-scripts/ifcfg-eth0:1
/etc/sysconfig/network-scripts/ifcfg-eth0:2

Multiple configuration files for a single eth0 interface;
this allows binding multiple IP addresses to a single NIC

/etc/sysconfig/network-scripts/route-eth0 (Red Hat)

Static route configuration for eth0

```
default 10.2.3.4 dev eth0
10.7.8.0/24 via 10.2.3.254 dev eth0
10.7.9.0/24 via 10.2.3.254 dev eth0
```

/etc/ethertypes (Red Hat)

Ethernet frame types; lists various Ethernet protocol
types used on Ethernet networks

/etc/network/interfaces (Debian)

List and configuration of all network interfaces

```
allow-hotplug eth0
iface eth0 inet static
    address 10.2.3.4
    netmask 255.255.255.0
    gateway 10.2.3.254
    dns-domain example.com
    dns-nameservers 8.8.8.8 4.4.4.4
```

/etc/hostname (Debian)

Hostname of the local machine

/etc/ethers (Debian)

ARP mappings

<code>dig example.org</code>	Perform a DNS lookup for the specified domain or hostname. Returns information in BIND zone file syntax; uses OS resolver libraries and hence does not use <code>/etc/resolv.conf</code>
<code>host example.org</code> <code>nslookup example.org</code>	Perform a DNS lookup for the specified domain or hostname. Does use <code>/etc/resolv.conf</code>
<code>dig @nameserver -t MX example.org</code> <code>host -t MX example.org nameserver</code>	Perform a DNS lookup for the MX record of the specified domain, querying <i>nameserver</i>
<code>dig example.org any</code> <code>host -a example.org</code>	Get all DNS records for a domain
<code>dig -x a.b.c.d</code> <code>host a.b.c.d</code>	Perform a reverse DNS lookup for the IP address <i>a.b.c.d</i>
<code>host -la example.org nameserver</code>	Perform a DNS Zone Transfer for zone <i>example.org</i> , querying the DNS server <i>nameserver</i> with a DNS ANY query. This lists all DNS records
<code>nslookup -norecurse example.org</code>	Check if the specified domain is present in the DNS cache
<code>whois example.org</code>	Query the WHOIS service for an Internet resource (usually a domain name)
<code>ping host</code>	Test if a remote host can be reached and measure the round-trip time to it. This is done by sending an ICMP Echo Request datagram and awaiting an ICMP Echo Response
<code>ping -M do -s size host</code>	Ping a remote host using an ICMP packet of size <i>size</i> (default is 56 bytes) and setting the DF (Don't Fragment) bit. This command can be used to find what is the maximum frame size allowed on the network, by trying increasingly higher values for <i>size</i> until it exceeds the MTU and the datagram is unable to reach the destination host
<code>fping -a host1 host2 host3</code>	Ping multiple hosts in parallel and report which ones are alive
<code>bing host1 host2</code>	Calculate point-to-point throughput between two hosts
<code>traceroute host</code>	Print the route, hop by hop, packets trace to a remote host. This is done by sending a sequence of ICMP Echo Request datagrams with increasing TTL values, starting with TTL=1, and expecting ICMP Time Exceeded datagrams
<code>tracpath host</code>	Simpler <code>traceroute</code>
<code>tcptraceroute host</code>	Implementation of <code>traceroute</code> that uses TCP packets
<code>mtr host</code>	<code>traceroute</code> and <code>ping</code> combined
<code>telnet host</code>	Establish a telnet connection to a remote host. Unencrypted and insecure, hence obsolete
<code>telnet host port</code>	Establish a telnet connection to a remote <i>host</i> on the specified <i>port</i> . Useful for a quick-and-dirty test of network services
<code>echo >/dev/tcp/ipaddress/port \ >/dev/null 2>&1 && echo "port is open"</code>	Check if <i>port</i> at <i>ipaddress</i> is open

<code>redir --laddr=ip1 --lport=port1 \</code> <code>--caddr=ip2 --cport=port2</code>	Redirect all connections, coming to local IP address <i>ip1</i> and port <i>port1</i> , to remote IP address <i>ip2</i> and port <i>port2</i>
<code>stunnel</code>	TLS encryption wrapper. Can be used to secure any client-server protocol
<code>lt</code>	Localtunnel. Exposes a web service on localhost to the Internet, via a unique public URL that proxies all web requests
<code>socat</code> <code>socat TCP-LISTEN:80,fork TCP:host:80</code> <code>socat TCP:timeserver:13 -</code>	Establish two bidirectional data stream and transfer data between them Forward local HTTP port to remote <i>host</i> 's HTTP port Query a <i>timeserver</i> using the Daytime Protocol
<code>wget</code> <code>wget --no-clobber --html-extension \</code> <code>--page-requisites --convert-links \</code> <code>--recursive --domains example.org \</code> <code>--no-parent www.example.org/path</code>	Download a file via HTTP, HTTPS, or FTP Download a whole website <i>www.example.org/path</i>
<code>curl</code> <code>curl www.example.org/file -o myfile</code> <code>curl -u user:psw 'ftp://server/file'</code> <code>curl -XPUT webserver -d'data'</code>	Transfer data to or from a remote host via HTTP, HTTPS, FTP, FTPS, SCP, SFTP, TFTP, DICT, TELNET, LDAP, or FILE Download a file via HTTP and save it locally under another name Download a file via FTP, after logging in to the server Send an HTTP PUT command with <i>data</i> to <i>webserver</i>
<code>aria2c</code>	Download a file via HTTP, HTTPS, FTP, BitTorrent, or Metalink
<code>tcpd</code>	Monitor and intercept incoming requests for services mapped one-to-one to executable files (e.g. telnet, finger, ftp, rsh, rlogin, tftp). <i>inetd</i> redirects these incoming service requests to <i>tcpd</i> , which logs the request and performs some checks before running the specific server program
<code>inetsim</code>	Simulate Internet services. This is useful when setting up a confined lab for malware analysis
<code>pktgen</code>	Network packet generator. Uses the DPDK (Data Plane Development Kit) packet processing framework
<code>trafgen</code>	Network packet generator
<code>packETH</code> <code>packETHcli</code>	Ethernet packet generator (GUI) Ethernet packet generator (command line)
<code>tllder</code>	Check the availability of a domain name across all valid TLDs
<code>uucp srchost!path desthost!path</code>	Unix-to-Unix copy. Copies files between hosts, identified by a bang path. Obsolete

<code>iwlist wlan0 scan</code>	List all wireless devices in range, with their quality of signal and other information
<code>iwlist wlan0 freq</code>	Display transmission frequency settings
<code>iwlist wlan0 rate</code>	Display transmission speed settings
<code>iwlist wlan0 txpower</code>	Display transmission power settings
<code>iwlist wlan0 key</code>	Display encryption settings
<code>iwgetid wlan0 option</code>	Print NWID, ESSID, AP/Cell address or other information about the wireless network that is currently in use
<code>iwconfig wlan0</code>	Display configuration of wireless interface <code>wlan0</code>
<code>iwconfig wlan0 option</code>	Configure wireless interface <code>wlan0</code>
<code>iw dev wlan0 station dump</code>	On a wireless card configured in AP Mode, display information (e.g. MAC address, tx/rx, bitrate, signal strength) about the clients
<code>rfkill list</code>	List installed wireless devices
<code>rfkill unblock n</code>	Enable wireless device number <i>n</i>
<code>hostapd</code>	Daemon that allows a wireless card to function in Host AP Mode, i.e. perform all functions of an Access Point
<code>hcidump -i device</code>	Display raw HCI (Host Controller Interface) data exchanged with a Bluetooth <i>device</i>

<code>netstat</code>	Display current network connections. Options: <code>-t</code> Display active TCP connections <code>-l</code> Display only listening sockets <code>-a</code> Display all listening and non-listening sockets <code>-p</code> Display PID and name of program to which each socket belongs <code>-i</code> Display network interfaces <code>-s</code> Display protocol statistics <code>-r</code> Display kernel routing tables (equivalent to <code>route -e</code>) <code>-n</code> Do not resolve hostnames or portnames <code>-c</code> Continuously display connections
<code>ss</code>	Display socket statistics (similarly to <code>netstat</code>)
<code>ss -t -a</code>	Display all TCP sockets
<code>arp-scan</code>	Scan all hosts on the current LAN. Uses ARP (Layer 2) packets; therefore it is able to also find hosts configured to drop all IP or ICMP traffic, and it cannot scan hosts outside the LAN
<code>snoop</code> (Solaris)	Packet sniffer
<code>tcptrace</code>	Tool for the analysis of TCP dump files such as those generated by <code>tcpdump</code> , <code>snoop</code> , etc.
<code>ipgrab</code>	Packet sniffer that includes full header fields
<code>dhcpcdump</code>	DHCP packet sniffer
<code>ngrep</code>	Filter data payload of network packets matching a specified regex
<code>nload</code>	Display a graph of the current network usage
<code>ntop</code> <code>ntopng</code>	Network usage analyzer
<code>nethogs</code>	Display bandwidth network usage by process
<code>iptraf</code> <code>iptraf-ng</code>	Interactive IP LAN monitor (ncurses UI)
<code>netserver</code>	Run a network performance benchmark server
<code>netperf</code>	Execute network performance benchmarks, by connecting to a <code>netserver</code> server
<code>iperf -s</code>	Run a network throughput benchmark server
<code>iperf -c server</code>	Execute network throughput tests in client mode, by connecting to an <code>iperf</code> server

In RHEL 7 and later, the network configuration is managed by the NetworkManager daemon.

A **connection** is a network configuration that applies to a **device** (aka network interface). A device can be included in multiple connections, but only one of them may be active at a time.

`/etc/sysconfig/network-scripts/ifcfg-connection`

Configuration for *connection* in ifcfg format. Deprecated

`/etc/NetworkManager/system-connections/`

Network profiles in keyfile format (INI file-based)

`nmcli`

NetworkManager CLI tool

`nmcli device status`

Show all network devices

`nmcli device disconnect iface`

Disconnects the device *iface*.

This command should be used instead of

`nmcli connection down connection`

because if *connection* is set to `autoconnect`, Network Manager will bring the connection (and the device) up again short time later

`nmcli connection show`

Show all connections.

Connections with an empty device entry are inactive

`nmcli connection show --active`

Show active connections

`nmcli connection show connection`

Show the configuration of *connection*

`nmcli connection add con-name connection \`
`type ethernet ifname iface ipv4.method manual \`
`ipv4.addresses 10.0.0.13/24 ipv4.gateway 10.0.0.254`

Configure a new *connection* that uses the Ethernet interface *iface* and assigns it an IPv4 address and gateway

`nmcli connection modify connection options`

Modify the configuration of *connection*

`nmcli connection up connection`

Brings up a *connection*

`nmcli connection reload`

Reload any manual change made to the files in
`/etc/NetworkManager/system-connections/`

`nmcli connection migrate`

Migrate all network profiles from ifcfg format to keyfile format

`man nmcli-examples`

Display examples of network configuration

`nmtui`

NetworkManager configuration tool (ncurses UI)

Network teaming allows binding together two or more network interfaces to increase throughput or provide redundancy. RHEL 7 and later implement network teaming via the `teamd` daemon.

How to set up a team connection

1. `nmcli connection add type team con-name teamcon ifname teamif \ config '{"runner":{"name":"loadbalance"}}'` Set up a team connection *teamcon* and a team interface *teamif* with a runner (in JSON code) for automatic failover
2. `nmcli connection modify teamcon ipv4.method manual \ ipv4.addresses 10.0.0.14/24 ipv4.gateway 10.0.0.254` Assign manually an IP address and gateway
3. `nmcli connection add type team-slave ifname iface \ master teamcon` Add an existing device *iface* as a slave of team *teamcon*. The slave connection will be automatically named `team-slave-iface`
4. Repeat the previous step for each slave interface.

<code>teamdctl <i>teamif</i> state</code>	Show the state of the team interface <i>teamif</i>
<code>teamnl <i>teamif</i> command</code>	Debug a team interface <i>teamif</i>
<code>man teamd.conf</code>	Display examples of team configurations and runners

A **network bridge** emulates a hardware bridge, i.e. a Layer 2 device able to forward traffic between networks based on MAC addresses.

How to set up a bridge connection

1. `nmcli connection add type bridge con-name brcon ifname brif` Set up a bridge connection *brcon* and a bridge interface *brif*
2. `nmcli connection modify brcon ipv4.method manual \ ipv4.addresses 10.0.0.15/24 ipv4.gateway 10.0.0.254` Assign manually an IP address and gateway
3. `nmcli connection add type bridge-slave ifname iface \ master brcon` Add an existing device *iface* as a slave of bridge *brcon*. The slave connection will be automatically named `bridge-slave-iface`
4. Repeat the previous step for each slave interface.

<code>brctl show <i>brif</i></code>	Show information about the bridge interface <i>brif</i>
<code>man nmcli-examples</code>	Display examples of teaming and bridging configuration

Network bonding aggregates physical or virtual network interfaces into a single logical one, in order to provide redundancy and achieve higher throughput.

How to set up a bond connection

- | | |
|--|---|
| <ol style="list-style-type: none">1. nmcli connection add type bond con-name <i>bondcon</i> ifname <i>bondif</i> \
bond.options "mode=active-backup"2. nmcli connection add type ethernet slave-type bond con-name \
<i>bondconport1</i> ifname <i>if1</i> master <i>bondcon</i>
nmcli connection add type ethernet slave-type bond con-name \
<i>bondconport2</i> ifname <i>if2</i> master <i>bondcon</i>3. nmcli connection modify <i>bondcon</i> ipv4.method manual \
ipv4.addresses 10.0.0.16/24 ipv4.gateway 10.0.0.2544. nmcli connection up <i>bondcon</i> | <p>Set up a bond connection <i>bondcon</i> and a bond interface <i>bondif</i> in active backup mode</p> <p>Create new connection profiles for the interfaces <i>if1</i> and <i>if2</i>, and assign them to the bond as slaves</p> <p>Assign manually an IP address and gateway</p> <p>Activate the connection</p> |
|--|---|

nmcli device show *bondcon*

Show the network bond status

cat /proc/net/bonding/*bondcon*

Show the slave interfaces of the network bond

Nmap is a network analyzer, auditing tool, and penetration testing tool. The GUI equivalent is **Zenmap**.

`nmap options host`

Scan a host, or all hosts in a subnet

Port state	
open	An application is listening for connections on the port
closed	No application is listening for connections on the port
filtered	Port is not responding to probe due to a firewall blocking the port, so port may be open or closed
unfiltered	Port is responding to probe, but it is impossible to tell whether port is open or closed

Scan technique options	
<code>-sT</code>	TCP connect scan aka full-open scan . Completes the three-way handshake; response will be SYN/ACK if port is open, RST if port is closed. Slow and likely to trigger IDS
<code>-sS</code>	SYN scan aka half-open scan or stealth scan . Sends a TCP packet with SYN flag set; same response as the TCP connect scan. Fast and stealth
<code>-sN</code>	Null scan . Sends a TCP packet with no flag set; response will be none if port is open, RST if port is closed. Stealth. Works only if the target host's OS TCP/IP implementation is based on RFC 793
<code>-sF</code>	FIN scan . Sends a TCP packet with FIN flag set. Same response as null scan. Scanning using a TCP packet with FIN and/or PSF and/or URG flags set is known as inverse TCP flag scan .
<code>-sX</code>	Xmas tree scan . Sends a TCP packet with FIN, PSF, and URG flags set. Same response as null scan
<code>-sA</code>	ACK scan . Sends a TCP packet with ACK flag set; response will be RST if port is open or closed (unfiltered), no response or ICMP error if port is filtered. Further analysis can be done on the TCP/IP RST response packet: if the TTL field is less than the max value, or the Window Size field is nonzero, port is open. Used to discover firewall rules, and to determine firewall type: if unfiltered ports are reported the firewall is stateless, otherwise the firewall is stateful
<code>-sI zombiehost</code>	IP ID idle scan . Uses a zombie host to perform the scan and predicts the port state of the target host by analyzing the IP fragmentation ID sequence numbers from the zombie host; if the sequence number has increased by 2 port is open, if it has increased by 1 port is closed. Fully stealth as no packets are sent from the scanning machine to the target host
<code>-sU</code>	UDP scan . Sends a UDP packet; response will be none if port is open or filtered, ICMP Port Unreachable (Type 3 Destination Unreachable) error if port is closed
<code>-sO</code>	IP protocol scan . Cycles through IP protocol numbers (instead of TCP or UDP ports) to try to determine which IP protocols (TCP, ICMP, IGMP, etc.) the target host supports; response will be any protocol response or none if port is open, ICMP Protocol Unreachable (Type 3 Destination Unreachable) error if port is closed, other ICMP Type 3 Destination Unreachable errors or none if port is filtered
<code>-sR</code>	RPC scan . Floods all TCP/UDP ports found open with SunRPC program NULL commands to try to determine whether they are RPC ports, and if yes, the service program and version number. It is recommended to use <code>-sV</code> instead as it gives more information
<code>-sC</code>	Script scan. Uses the default script set. The Nmap Scripting Engine permits writing scripts (in the Lua programming language) to perform automatically various types of network scans
<code>--script=name</code>	Script scan. Runs the <i>name</i> scan script. Examples: <code>--script=sniffer-detect host</code> Check if a host has its NIC in promiscuous mode (sniffer) <code>--script=firewalk host</code> Attempt to detect firewall or gateway rules <code>--script=http-trace -d host</code> Send an HTTP TRACE request to find if TRACE method is enabled <code>--script=http-enum host</code> Enumerate dirs used by common web applications and web servers

Host discovery options	
-sL	List scan. Simple network host discovery, with reverse DNS resolution. No packet is sent to target hosts
-sn	No port scan aka ping sweep. Used to detect how many hosts are up
-Pn	No ping; skip host discovery. The subsequent scan operation will be performed against all hosts, instead of only those discovered to be up during this phase
-PSport	Send a TCP SYN packet to the specified port
-PAport	Send a TCP ACK packet to the specified port
-PUport	Send a UDP packet to the specified port
-PYport	Send a SCTP packet containing a minimal INIT chunk to the specified port
-PE	Send a ICMP Echo Request (type 8)
-PP	Send a ICMP Timestamp Request (type 13)
-PM	Send a ICMP Address Mask Request (type 17)
-POprotocol	Send IP packets with the specified protocol number set in their header
-PR	ARP scan. Default discovery type when scanning the current LAN
--traceroute	After the scan, trace path to host to determine port and protocol most likely to reach the target host
IDS evasion, firewall evasion, and spoofing options	
-f	Use tiny fragmented packets (8 bytes or less) for IDS evasion. Might crash the target host
--mtu offset	Use fragmented packets of size <i>offset</i> (must be a multiple of 8) for IDS evasion
-D ip	Spoof the scanning machine IP address as <i>ip</i>
-D RND:n	Spoof the scanning machine IP address using <i>n</i> randomly generated addresses. The real IP address is included among the decoys
--ip-options "L ip1 ip2"	Use loose source routing for IDS evasion, requiring that the packet is loose source routed through the waypoints with IP address <i>ip1</i> and <i>ip2</i>
--ip-options "S ip1 ip2"	Use strict source routing for IDS evasion, requiring that the packet is strictly source routed through the waypoints with IP address <i>ip1</i> and <i>ip2</i> . All waypoints must be specified
Timing options	
-T0	Paranoid. Extremely slow serialized scan for IDS evasion. Will take a long time to complete
-T1	Sneaky. Very slow serialized scan for IDS evasion. Will take a long time to complete
-T2	Polite. Slow serialized scan to consume less bandwidth and resources of the target
-T3	Normal. Parallel scan. Default
-T4	Aggressive. Fast parallel scan, to be used on networks with a high bandwidth. Recommended
-T5	Insane. Very fast parallel scan, to be used on networks with a very high bandwidth. Might be less accurate

Other options	
-A	Aggressive scan. Equivalent to -O -sV -sC --traceroute
-O	OS fingerprinting, to find out which operating system is running on target host
-sV	Version detection, to determine protocol, application name, version number, device type, etc.
-6	Enable IPv6 scanning
-p <i>port</i>	Scan only the specified port or port range, instead of the most common 1000 ports for each protocol
-p-	Scan all ports (from 1 to 65535)
--top-ports <i>n</i>	Scan only the <i>n</i> most popular ports
-F	Fast mode; scan fewer ports than the default, hence enumerating all hosts faster
-r	Scan ports in numerical order, instead of random order
-n	Do not do DNS resolution
-R	Always do DNS resolution
-oN <i>file.nmap</i>	Save output to <i>file</i> in standard format (slightly different from interactive mode output)
-oX <i>file.xml</i>	Save output to <i>file</i> in XML format

Tcpdump is a packet sniffer (aka packet analyzer) which uses the `libpcap` library for packet capture. The GUI equivalent of tcpdump is **Wireshark**, formerly called Ethereal.

Sniffers operate at the Data Link layer (Layer 2).

On a wired medium, for a sniffing machine to be able to capture all network traffic, and not only the traffic from/to the machine itself, the machine's NIC must be set to promiscuous mode. Furthermore, only traffic within a network segment connected via a hub (i.e. the collision domain) can be sniffed; in the case of a switched network, the sniffing machine needs to be connected to the switch's SPAN port (which performs port mirroring) in order to be able to capture all traffic. In the case of a wireless NIC, the chipset also determines capabilities for modes of operation.

Active sniffing refers to sniffing through a switch. Passive sniffing refers to sniffing through a hub.

`tcpdump options expression`

Print the content of sniffed packets that match *expression*.

Options:

- `-v -vv` Increasing levels of verbosity
- `-n` Do not perform DNS resolution on host addresses
- `-nn` Do not convert protocol and port numbers to names

`tcpdump -i eth0`

Sniff all network traffic on interface `eth0`

`tcpdump ip host 10.0.0.2 tcp port 25`

Sniff network packets on TCP port 25 from and to 10.0.0.2

`tcpdump ether host '45:67:89:ab:cd:ef'`

Sniff traffic from and to the network interface having MAC address 45:67:89:ab:cd:ef

`tcpdump 'src host 10.0.0.2 and \`
`(tcp port 80 or tcp port 443)'`

Sniff HTTP and HTTPS traffic having as source host 10.0.0.2

`tcpdump -i eth0 not port 22`

Sniff all traffic on `eth0` except that belonging to a SSH connection

`tcpdump -i eth0 arp`

Sniff ARP traffic on `eth0`

`tcpdump ip host 10.0.0.2 and not 10.0.0.9`

Sniff IP traffic between 10.0.0.2 and any other host except 10.0.0.9

PCAP filter syntax

<code>tcp.port==25 or icmp</code>	Show SMTP and ICMP traffic		
<code>ip.addr==10.0.0.2</code>	Show traffic from and to 10.0.0.2		
<code>ip.src==10.0.0.2 or ip.dst==10.0.0.2</code>			
<code>ip.src==10.0.0.3 and frame.pkt_len > 400</code>	Show packets coming from 10.0.0.3 with frame length higher than 400		
<code>http.request</code>	Show HTTP requests		
<code>udp contains 76:54</code>	Show UDP packets containing the 2-byte hex sequence 0x76, 0x54 in the header or the payload, at any offset		

<code>==</code> <code>eq</code>	Equal to	<code>></code> <code>gt</code>	Greater than	<code>>=</code> <code>ge</code>	Greater than or equal to	<code>&&</code> <code>and</code>	Logical AND
<code>!=</code> <code>ne</code>	Not equal to	<code><</code> <code>lt</code>	Less than	<code><=</code> <code>le</code>	Less than or equal to	<code> </code> <code>or</code>	Logical OR

Netcat is "the Swiss Army knife of networking", a very flexible generic TCP/IP client/server.

<pre>nc -z 10.0.0.7 22 ncat 10.0.0.7 22 (RHEL) netcat 10.0.0.7 22 (SUSE)</pre>	Scan for a listening SSH daemon on remote host 10.0.0.7
<pre>nc -l -p 25</pre>	Listen for connections on port 25 (i.e. mimic an SMTP server). Send any input received on stdin to the connected client and dump on stdout any data received from the client
<pre>nc 10.0.0.7 389 < file</pre>	Push the content of <i>file</i> to port 389 on remote host 10.0.0.7
<pre>echo "GET / HTTP/1.0\r\n\r\n" nc 10.0.0.7 80</pre>	Connect to web server 10.0.0.7 and issue an HTTP GET
<pre>while true; \ do nc -l -p 80 -q 1 < page.html; done while true; \ do echo "<html><body>Hello</body></html>" \ ncat -l -p 80; done</pre>	Start a minimal web server, serving the specified HTML page to clients
<pre>nc -v -n -z -w1 -r 10.0.0.7 1-1023</pre>	Run a TCP port scan against remote host 10.0.0.7. Probes randomly all privileged ports with a 1-second timeout, without resolving service names, and with verbose output
<pre>echo "" nc -v -n -w1 10.0.0.7 1-1023</pre>	Retrieve the greeting banner of any network service that might be running on remote host 10.0.0.7

Hping3 is a packet crafting tool, able to send any custom TCP/IP packet to a remote host and display the reply. It is an extension of hping2, and is command-line compatible with it while having extended capabilities for packet generation.

`hping3 options host` Send a crafted packet to *host*. By default, it sends TCP headers to port 0 of remote host with no TCP flag set and a window size of 64

hping3 options	
<code>-c n</code>	Send <i>n</i> packets
<code>-p n</code>	Use port <i>n</i>
<code>-a src</code> <code>--spoof src</code>	Set <i>src</i> as a fake IP source address for sent packets
<code>-1</code> <code>--ICMP</code>	Use the ICMP protocol. By default, hping3 uses TCP
<code>-2</code> <code>--UDP</code>	Use the UDP protocol
<code>-8 n1-n2</code> <code>--scan n1-n2</code>	Operate in scan mode, scanning the port range from <i>n1</i> to <i>n2</i>
<code>-9 signature</code> <code>--listen signature</code>	Operate in listening mode, trying to intercept <i>signature</i>
<code>-A</code>	Set the ACK flag in probe packets. An ACK scan can be used to check if the remote host is alive, when it does not respond to ping packets
<code>-S</code>	Set the SYN flag in probe packets
<code>-F</code>	Set the FIN flag in probe packets
<code>-P</code>	Set the PSH flag in probe packets
<code>-U</code>	Set the URG flag in probe packets
<code>-Q</code>	Collect all TCP sequence numbers generated by the remote host
<code>--tcp-timestamp</code>	Attempt to guess the timestamp update frequency and uptime of the remote host

`hping3 -S -p 25 -c 5 host` Send 5 TCP packets, with the SYN flag set, to port 25 of remote host

`hping3 --scan 1-1024 -S host` Perform a SYN scan on ports 1 to 1024 against the remote host

`hping3 --udp --rand-source --data 512 host` Send UDP packets with random source address and a data body size of 512 bytes

`hping3 -S -p 80 --flood host` Perform a TCP SYN flood DoS attack against a webserver

`hping3 -A -p 25 host` Verify if a mailserver is alive (if it is, it will reply with an RST)

The **TCP Wrapper** feature provides basic traffic filtering of incoming network connections. To use this feature, the service binary must have been compiled with the `libwrap.a` library.

```
ldd service_binary | grep libwrap
```

Find if a network service is TCP Wrapped

`/etc/hosts.allow`
`/etc/hosts.deny`

Host access control files used by the TCP Wrapper system.

Each file contains zero or more *daemon:client* lines. The first matching line is considered.

Access is granted when a *daemon:client* pair matches an entry in `/etc/hosts.allow`.
Otherwise, access is denied when a *daemon:client* pair matches an entry in `/etc/hosts.deny`.
Otherwise, access is granted.

/etc/hosts.allow and /etc/hosts.deny lines syntax	
ALL: ALL	All services to all hosts
ALL: .example.edu	All services to all hosts of the example.edu domain
ALL: .example.edu EXCEPT host1.example.edu	All services to all hosts of example.edu, except host1
in.fingerd: .example.com	Finger service to all hosts of example.com
in.tftpd: LOCAL	TFTP to hosts of the local domain only
sshd: 10.0.0.3 10.0.0.4 10.1.1.0/24	SSH to the hosts and network specified
sshd: 10.0.1.0/24 sshd: 10.0.1. sshd: 10.0.1.0/255.255.255.0	SSH to 10.0.1.0/24
in.tftpd: ALL: spawn (/safe_dir/safe_finger -l @%h \ /bin/mail -s %d-%h root) &	Send a finger probe to hosts attempting TFTP and notify the root user via email
portmap: ALL: (echo Illegal RPC request from %h \ /bin/mail root) &	When a client attempts an RPC request via portmapper (NFS access), echo a message to the terminal and notify the root user via email

Output of command <code>route -en</code>							
Kernel IP routing table							
Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
192.168.3.0	0.0.0.0	255.255.255.0	U	0	0	0	eth0
0.0.0.0	192.168.3.1	0.0.0.0	UG	0	0	0	eth0
Destination	<i>network or host</i>	destination network or host					
	0.0.0.0	default route					
Gateway	<i>host</i>	gateway					
	0.0.0.0						
	*	no gateway needed, network is directly connected					
Genmask	-	rejected route					
	<i>network mask</i>	network mask to apply for the destination network					
	255.255.255.255	destination host					
Flags	0.0.0.0	default route					
	U	route is up					
	G	use gateway					
	H	target is host					
	!	rejected route					
	D	dynamically installed by daemon					
	M	modified from routing daemon					
	R	reinstate route for dynamic routing					

```
ip route
route -en
route -F
netstat -rn
```

Display IP routing table

```
ip route show cache
route -C
```

Display kernel routing cache

```
ip route add default via 10.1.1.254
route add default gw 10.1.1.254
```

Add a default gateway 10.1.1.254

```
ip route add 10.2.0.1 dev eth0
ip route add 10.2.0.1 via 10.2.0.254
route add -host 10.2.0.1 gw 10.2.0.254
```

Add a route for a host 10.2.0.1

```
ip route add 10.2.0.0/16 via 10.2.0.254
route add -net 10.2.0.0 netmask 255.255.0.0 gw 10.2.0.254
```

Add a route for a network 10.2.0.0/16

```
ip route delete 10.2.0.1 dev eth0
route del -host 10.2.0.1 gw 10.2.0.254
```

Delete a route for a host 10.2.0.1

```
ip route flush all
```

Delete the routing table for all interfaces

```
rdisc
```

Network router discovery daemon.
Client for IRDP (ICMP Router Discover Protocol).
Runs at bootup to populate the network routing tables with default routes

The Netfilter framework provides firewalling (packet filtering and NAT routing) capabilities in Linux. It is implemented by **ipfirewall** in the 2.0 kernel, **ipchains** in the 2.2 kernel, **iptables** in the 2.4 kernel, and **nftables** in the 3.13 kernel; each one of these is managed via the appropriate user-space application program. The ability to track connection state is provided by the `ip_conntrack` kernel module.

In RHEL 6, the service `iptables` provides all firewall functionalities.
In RHEL 7, the firewall is managed by the `firewalld` daemon, which uses `iptables` as backend. It is possible, but discouraged, to use `iptables` directly by disabling `firewalld` and installing the package `iptables-services`; this provides `systemd` units for `iptables`, which is implemented in the kernel and therefore does not have a daemon process or a service. In RHEL 8 and 9, the firewall is managed by `firewalld`, with `nftables` as backend.
In Ubuntu, firewall capabilities are provided by the `ufw` (Uncomplicated Firewall) service, with `iptables` as backend.

<code>ipfwadm</code>	Administration tool for <code>ipfirewall</code> . Obsolete
<code>ipchains</code>	Administration tool for <code>ipchains</code> . Obsolete
<code>iptables</code>	Administration tool for <code>iptables</code> on IPv4
<code>ip6tables</code>	Administration tool for <code>iptables</code> on IPv6
<code>system-config-firewall</code>	GUI frontend for <code>iptables</code>
<code>iptables -L</code>	List all rules in all chains. By default, rules are stored in <code>/etc/sysconfig/iptables</code>
<code>iptables-restore < file</code>	Load into <code>iptables</code> a dump of firewall rules specified in <i>file</i>
<code>iptables-save > file</code>	Dump the current <code>iptables</code> firewall rules to <i>file</i>


iptables rules dump file	
*filter :INPUT ACCEPT [0:0] :FORWARD ACCEPT [0:0] :OUTPUT ACCEPT [0:0] COMMIT	Delete all rules and open the firewall to all connections
*filter :INPUT DROP [45:2307] :FORWARD ACCEPT [0:0] :OUTPUT ACCEPT [1571:4260654] -A INPUT -m state --state RELATED,ESTABLISHED -j ACCEPT -A INPUT -p icmp -j ACCEPT -A INPUT -i lo -j ACCEPT -A INPUT -p tcp -m state --state NEW -m tcp --dport 22 -j ACCEPT COMMIT	Open the firewall to SSH connections only.

The `[packet_counter:byte_counter]` values are for debug purposes only.

<code>nft</code>	Administration tool for <code>nftables</code>
<code>ufw</code>	Administration tool for Uncomplicated Firewall

Iptables uses **tables** containing sets of **chains**, which contain sets of **rules**. Each rule has a **target** (e.g. ACCEPT). The "filter" table contains chains INPUT, FORWARD, OUTPUT (built-in chains); this is the default table to which all iptables commands are applied, unless another table is specified via the `-t` option. The "nat" table contains chains PREROUTING, OUTPUT, POSTROUTING. The "mangle" table contains chains PREROUTING, OUTPUT. When a packet enters the system, it is handed to the INPUT chain. If the destination is local, it is processed; if the destination is not local and IP forwarding is enabled, the packet is handed to the FORWARD chain, otherwise it is dropped. An outgoing packet generated by the system will go through the OUTPUT chain. If NAT is in use, an incoming packet will pass at first through the PREROUTING chain, and an outgoing packet will pass last through the POSTROUTING chain.

<code>iptables -A INPUT -s 10.0.0.6 -j ACCEPT</code>	Add a rule to accept all packets from 10.0.0.6
<code>iptables -A INPUT -s 10.0.0.7 -j REJECT</code>	Add a rule to reject all packets from 10.0.0.7 and send back a ICMP response to the sender
<code>iptables -A INPUT -s 10.0.0.8 -j DROP</code>	Add a rule to silently drop all packets from 10.0.0.8
<code>iptables -A INPUT -s 10.0.0.9 -j LOG</code>	Add a rule to log (via syslog) all packets from 10.0.0.9
<code>iptables -D INPUT -s 10.0.0.9 -j LOG</code>	Delete a specific rule
<code>iptables -D INPUT 42</code>	Delete rule 42 of the INPUT chain
<code>iptables -F INPUT</code>	Flush all rules of the INPUT chain
<code>iptables -F</code>	Flush all rules, hence disabling the firewall
<code>iptables -t mangle -F</code>	Flush all rules of the "mangle" table
<code>iptables -t mangle -X</code>	Delete all user-defined (not built-in) rules in the "mangle" table
<code>iptables -L INPUT</code>	List the rules of the INPUT chain
<code>iptables -L -n</code>	List all rules, without translating numeric values (IP addresses to FQDNs and port numbers to services)
<code>iptables -N mychain</code>	Define a new chain
<code>iptables -P INPUT DROP</code>	Define the chain policy target, which takes effect when no rule matches and the end of the rules list is reached
<code>iptables -A OUTPUT -d 10.7.7.0/24 -j DROP</code>	Add a rule to drop all packets with destination 10.7.7.0/24
<code>iptables -A FORWARD -i eth0 -o eth1 -j LOG</code>	Add a rule to log all packets entering the system via eth0 and exiting via eth1
<code>iptables -A INPUT -p 17 -j DROP</code> <code>iptables -A INPUT -p udp -j DROP</code>	Add a rule to drop all incoming UDP traffic (protocol numbers are defined in <code>/etc/protocols</code>)
<code>iptables -A INPUT --sport 1024:65535 --dport 53 \</code> <code>-j ACCEPT</code>	Add a rule to accept all packets coming from any unprivileged port and with destination port 53
<code>iptables -A INPUT -p icmp --icmp-type echo-request \</code> <code>-m limit --limit 1/s -i eth0 -j ACCEPT</code>	Add a rule to accept incoming pings through eth0 at a maximum rate of 1 ping/second
<code>iptables -A INPUT -m state --state ESTABLISHED \</code> <code>-j ACCEPT</code>	Load the module for stateful packet filtering, and add a rule to accept all packets that are part of a communication already tracked by the state module
<code>iptables -A INPUT -m state --state NEW -j ACCEPT</code>	Add a rule to accept all packets that are not part of a communication already tracked by the state module
<code>iptables -A INPUT -m state --state RELATED -j ACCEPT</code>	Add a rule to accept all packets that are related (e.g. ICMP responses to TCP or UDP traffic) to a communication already tracked by the state module
<code>iptables -A INPUT -m state --state INVALID -j ACCEPT</code>	Add a rule to accept all packets that do not match any of the states above



SNAT (Source Network Address Translation)

```
iptables -t nat -A POSTROUTING -s 10.0.0.0/24 -o eth1 \
-j SNAT --to-source 93.184.216.119
```

Map all traffic leaving the LAN to the external IP address 93.184.216.119

```
iptables -t nat -A POSTROUTING -s 10.0.0.0/24 -o eth1 \
-j SNAT --to-source 93.184.216.119:93.184.216.127
```

Map all traffic leaving the LAN to a pool of external IP addresses 93.184.216.119-127

```
iptables -t nat -A POSTROUTING -o eth1 -j MASQUERADE
```

Map all traffic leaving the LAN to the address dynamically assigned to eth1 via DHCP

DNAT (Destination Network Address Translation)

```
iptables -t nat -A PREROUTING -i eth1 -d 93.184.216.119 \
-j DNAT --to-destination 10.0.0.13
```

Allow the internal host 10.0.0.13 to be publicly reachable via the external address 93.184.216.119

PAT (Port Address Translation)

```
iptables -t nat -A PREROUTING -i eth1 -d 93.184.216.119 \
-p tcp --dport 80 -j DNAT --to-destination 10.0.0.13:8080
```

Make publicly accessible a webserver that is located in the LAN, by mapping port 8080 of the internal host 10.0.0.13 to port 80 of the external address 93.184.216.119

```
iptables -t nat -A PREROUTING -i eth0 -d ! 10.0.0.0/24 \
-p tcp --dport 80 -j REDIRECT --to-ports 3128
```

Redirect all outbound HTTP traffic originating from the LAN to a proxy running on port 3128 on the Linux box

```
sysctl -w net.ipv4.ip_forward=1
echo 1 > /proc/sys/net/ipv4/ip_forward
```

Enable IP forwarding; necessary to set up a Linux machine as a router. (This command causes other network options to be changed as well.)

In firewalld, a network interface (aka **interface**) or a subnet address (aka **source**) can be assigned to a specific **zone**. To determine to which zone a packet belongs, first the zone of the source is analyzed, then the zone of the interface; if no source or interface matches, the packet is associated to the default zone (which is "public", unless set otherwise).

A command is applied to the default zone, unless a different zone is specified (via `--zone=zone`).

By default, commands are temporary; adding the `--permanent` option to a command sets it as permanent, or shows permanent settings only.

Temporary commands are effective immediately but are canceled at reboot, firewall reload, or firewall restart.

Permanent commands are effective only after reboot, firewall reload, or firewall restart.

Firewalld zones	
block	Rejects incoming connections with an ICMP HOST_PROHIBITED; allows only established connections
dmz	Used to expose services to the public; allows only specific incoming connections
drop	Drops all incoming packets; allows only outgoing connections
external	Used for routing and masquerading; allows only specific connections
home	Allows only specific incoming connections
internal	Used to define internal networks and allow only private network traffic
public	Allows only specific incoming connections. Default zone
trusted	Accepts all traffic
work	Used to define internal networks and allow only private network traffic

The list of firewall zones can be obtained via the command `firewall-cmd --get-zones`.

```
systemctl status firewalld
firewall-cmd --state
```

Check the status of the firewall

```
firewall-config
```

Firewall management GUI

```
firewall-cmd --reload
```

Reload firewall configuration; this applies all permanent changes and cancels all temporary changes. Current connections are not terminated

```
firewall-cmd --complete-reload
```

Reload firewall configuration, stopping all current connections

```
firewall-cmd --runtime-to-permanent
```

Transform all temporary changes to permanent

```
firewall-cmd --list-all-zones
```

List all zones and their full settings

```
firewall-cmd --get-default-zone
```

Show the default zone

```
firewall-cmd --set-default-zone=home
```

Set "home" as the default zone

```
firewall-cmd --get-active-zones
```

Show the active zones i.e. zones bound to either an interface or a source

```
firewall-cmd --get-zones
```

Show all available zones

```
firewall-cmd --get-zone-of-interface=eth0
```

Show the zone assigned to eth0

```
firewall-cmd --new-zone=test
```

Create a new zone called "test"

```
firewall-cmd --zone=home --change-interface=eth0
```

Assign eth0 to the "home" zone

```
firewall-cmd --zone=home --add-source=10.1.1.0/24
```

Assign 10.1.1.0/24 to the "home" zone i.e. route all traffic from that subnet to that zone

```
firewall-cmd --zone=home --list-sources
```

List sources bound to the "home" zone

```
firewall-cmd --list-all
```

List temporary settings

```
firewall-cmd --list-all --permanent
```

List permanent settings

<code>firewall-cmd --add-service=ssh</code>	Add the SSH service to the list of allowed services
<code>firewall-cmd --add-port=22/tcp</code>	
<code>firewall-cmd --add-service={ssh,http,https}</code>	Allow SSH, HTTP, and HTTPS services
<code>firewall-cmd --list-services</code>	Show allowed services
<code>firewall-cmd --list-ports</code>	Show open ports
<code>firewall-cmd --get-services</code>	List all predefined services
Predefined services are configured in <code>/usr/lib/firewalld/services/service.xml</code> .	
User-defined services are configured in <code>/etc/firewalld/services/service.xml</code> .	
<code>firewall-cmd --get-icmptypes</code>	Show all known types of ICMP messages
<code>firewall-cmd --add-icmp-block=echo-reply</code>	Block a specific ICMP message type
<code>firewall-cmd --query-icmp-block=echo-reply</code>	Tell if a specific ICMP message type is blocked
<code>firewall-cmd --list-icmp-block</code>	Show the list of blocked ICMP message types
<code>firewall-cmd --add-rich-rule='richrule'</code>	Set up a rich rule (used for more complex and detailed firewall configurations)
<code>firewall-cmd --add-rich-rule='rule \</code> <code>family=ipv4 source address=10.2.2.0/24 service name=tftp</code> <code>log prefix=tftp level=info limit value=3/m accept'</code>	Set up a rich rule to allow TFTP connections from subnet 10.2.2.0/24 and log them via syslog at a rate of 3 per minute
<code>firewall-cmd --list-rich-rules</code>	List all rich rules
The manpage <code>man firewalld.richlanguage</code> contains examples of rich rules.	
<code>firewall-cmd --direct --add-rule directrule</code>	Set up a direct rule (in iptables format)
<code>firewall-cmd --direct --add-rule \</code> <code>ipv4 filter INPUT 0 -p tcp --dport 22 -j ACCEPT</code>	Set up a direct rule to allow SSH connections
<code>firewall-offline-cmd directrule</code>	Set up a direct rule when firewalld is not running
<code>firewall-cmd --direct --get-all-rules</code>	Show all direct rules
User-defined direct rules are stored in <code>/etc/firewalld/direct.xml</code> .	
The manpage <code>man firewalld.direct</code> documents the syntax of direct rules.	
<code>firewall-cmd --zone=zone --add-masquerade</code>	Set up masquerading for hosts of <i>zone</i> ; packets originating from <i>zone</i> will get the firewall's IP address on the "external" zone as source address
<code>firewall-cmd --zone=zone --add-rich-rule='rule \</code> <code>family=ipv4 source address=10.2.2.0/24 masquerade'</code>	Set up masquerading only for those hosts of <i>zone</i> located in subnet 10.2.2.0/24
<code>firewall-cmd --zone=zone --add-forward-port=\</code> <code>port=22:proto=tcp:toport=2222:toaddr=10.7.7.7</code>	Set up port forwarding for hosts of <i>zone</i> ; incoming connections to port 22 for hosts of <i>zone</i> will be forwarded to port 2222 on host 10.7.7.7

Secure Shell (SSH) is a protocol (not a shell) for encrypted secure communications. It is mostly used as a replacement to Telnet to securely login to a remote server's terminal, but can be applied to any network protocol: some of the most common applications of SSH are **SCP (Secure Copy)** and **SFTP (Secure File Transfer Protocol aka SSH File Transfer Protocol)**.

<code>ssh user@host</code>	Connect to a remote <i>host</i> via SSH and login as <i>user</i> . Options: -v -vv -vvv Increasing levels of verbosity -p <i>n</i> Use port <i>n</i> instead of standard port 22
<code>ssh user@host command</code>	Execute a command on a remote host
<code>autossh user@host</code>	Connect to a remote host, monitoring the connection and restarting it automatically if it dies
<code>sshpass -p password ssh user@host</code>	Connect to a remote host using the specified password
<code>pssh -i -H "host1 host2 host3" command</code>	Execute a command in parallel on a group of remote hosts
<code>ssh-keygen -t rsa -b 2048</code>	Generate interactively a 2048-bit RSA key pair; will prompt for a passphrase
<code>ssh-keygen -t dsa</code>	Generate a DSA key pair
<code>ssh-keygen -p -t rsa</code>	Change passphrase of the private key
<code>ssh-keygen -q -t rsa -f keyfile -N '' -C ''</code>	Generate an RSA key with no passphrase (for non-interactive use) and no comment
<code>ssh-keygen -lf keyfile</code>	View key length and fingerprint of a public or private key
<code>< keyfile.pub awk '{print \$2}' \</code> <code> base64 -d openssl hashfunction</code>	View fingerprint of a key, calculated using <i>hashfunction</i> . RSA keys fingerprint use <i>sha1</i> (deprecated) or <i>md5</i>
<code>ssh-keyscan host >> ~/.ssh/known_hosts</code>	Get the public key of <i>host</i> and add it to the user's known hosts file
<code>ssh-agent</code>	Echo to the terminal the environment variables that must be set in order to use the SSH Agent
<code>eval `ssh-agent`</code>	Start the SSH Agent daemon that caches decrypted private keys in memory; also shows the PID of ssh-agent and sets the appropriate environment variables. Once ssh-agent is started, the keys to cache must be added via the <code>ssh-add</code> command; cached keys will then be automatically used by any SSH tool e.g. <code>ssh</code> , <code>sftp</code> , <code>scp</code>
<code>ssh-agent bash -c 'ssh-add keyfile'</code>	Start ssh-agent and cache the specified key
<code>ssh-add</code>	Add the default private keys to the ssh-agent cache
<code>ssh-add keyfile</code>	Add a specific private key to the ssh-agent cache
<code>ssh-copy-id user@host</code>	Use locally available keys to authorize, via public key authentication, login of <i>user</i> on a remote <i>host</i> . This is done by copying the user's local public key <code>~/.ssh/id_rsa.pub</code> to <code>~/.ssh/authorized_keys</code> on the remote host

```
scp /path1/file user@host:/path2/  
scp user@host:/path1/file /path2/  
scp user1@host1:/path1/file user2@host2:/path2/
```

Non-interactive secure file copy via SSH.
Can transfer files from local to remote, from remote to local,
or between two remote hosts

```
sftp user@host
```

SSH FTP-like tool for secure file transfer

```
scponly
```

SSH wrapper pseudo-shell providing access to remote users
for secure file transfer, but without execution privileges

```
sshfs user@host:/dir/ mountpoint/
```

SSH tool that allows mounting a remote directory as an SSH
filesystem on a mountpoint on the local machine. Uses the
FUSE kernel module.
The filesystem can be unmounted via the command
`fusermount -u mountpoint/`

SSH port forwarding (aka SSH tunneling)

```
ssh -L 2525:mail.foo.com:25 user@mail.foo.com
```

Establish a SSH encrypted tunnel from localhost to remote host mail.foo.com, redirecting traffic from local port 2525 to port 25 of remote host mail.foo.com.

Useful if the local firewall blocks outgoing port 25. In this case, port 2525 is used to go out; the application must be configured to connect to localhost on port 2525 (instead of mail.foo.com on port 25)

```
ssh -L 2525:mail.foo.com:25 user@login.foo.com
```

Establish a SSH encrypted tunnel from localhost to remote host login.foo.com.

Remote host login.foo.com will then forward, unencrypted, all data received over the tunnel on port 2525 to remote host mail.foo.com on port 25

SSH reverse forwarding (aka SSH reverse tunneling)

```
ssh -R 2222:localhost:22 user@login.foo.com
```

Establish a SSH encrypted reverse tunnel from remote host login.foo.com back to localhost, redirecting traffic sent to port 2222 of remote host login.foo.com back towards local port 22.

Useful if the local firewall blocks incoming connections so remote hosts cannot connect back to local machine. In this case, port 2222 of login.foo.com is opened for listening and connecting back to localhost on port 22; remote host login.foo.com is then able to connect to the local machine on port 2222 (redirected to local port 22)

SSH as a SOCKS proxy

```
ssh -D 33333 user@login.foo.com
```

The application supporting SOCKS must be configured to connect to localhost on port 33333. Data is tunneled from localhost to login.foo.com, then unencrypted to destination

X11 Forwarding

```
ssh -X user@login.foo.com
```

Enable the local display to execute locally an X application stored on a remote host login.foo.com

How to enable public key authentication

1. On remote host, set `PubkeyAuthentication yes` in `/etc/ssh/sshd_config`
2. On local machine, do `ssh-copy-id you@remotehost` (or copy your public key to the remote host by hand)

How to enable host-based authentication amongst a group of trusted hosts

1. On all hosts, set `HostbasedAuthentication yes` in `/etc/ssh/sshd_config`
2. On all hosts, create `/etc/ssh/shosts.equiv` and enter in this file all trusted hostnames
3. Connect via SSH manually from your machine on each host so that all hosts' public keys go into `~/.ssh/known_hosts`
4. Copy `~/.ssh/known_hosts` from your machine to `/etc/ssh/ssh_known_hosts` on all hosts

How to enable X11 Forwarding

1. On remote host 10.2.2.2, set `X11Forwarding yes` in `/etc/ssh/sshd_config`, and make sure that `xauth` is installed
2. On local host 10.1.1.1, type `ssh -X 10.2.2.2`, then run on remote host the graphical application e.g. `xclock &`

It is also possible to enable X11 Forwarding via telnet (unencrypted, therefore insecure and not recommended):

1. On remote host 10.2.2.2, type `export DISPLAY=10.1.1.1:0.0`
2. On local host 10.1.1.1, type `xhost +`
3. On local host 10.1.1.1, type `telnet 10.2.2.2`, then run on remote host the graphical application e.g. `xclock &`

<code>/etc/ssh/sshd_config</code>	SSH server daemon configuration file
<code>/etc/ssh/ssh_config</code>	SSH client global configuration file
<code>/etc/ssh/ssh_host_key</code>	Host's private key (should be mode 0600)
<code>/etc/ssh/ssh_host_key.pub</code>	Host's public key
<code>/etc/ssh/shosts.equiv</code>	Names of trusted hosts for host-based authentication
<code>/etc/ssh/ssh_known_hosts</code>	Database of host public keys that were previously accepted as legitimate
<code>~/.ssh/</code>	User's SSH directory (must be mode 0700)
<code>~/.ssh/config</code>	SSH client user configuration file
<code>~/.ssh/id_rsa</code> <code>~/.ssh/id_dsa</code>	User's RSA or DSA private key, as generated by <code>ssh-keygen</code>
<code>~/.ssh/id_rsa.pub</code> <code>~/.ssh/id_dsa.pub</code>	User's RSA or DSA public key, as generated by <code>ssh-keygen</code>
<code>~/.ssh/known_hosts</code>	Host public keys that were previously accepted as legitimate by the user
<code>~/.ssh/authorized_keys</code> <code>~/.ssh/authorized_keys2</code> (obsolete)	Trusted public keys; the corresponding private keys allow the user to authenticate on this host

<code>/etc/ssh/sshd_config</code>	SSH server configuration file
<code>PermitRootLogin yes</code>	Control superuser login via SSH. Possible values are: <code>yes</code> Superuser can login <code>no</code> Superuser cannot login <code>without-password</code> Superuser cannot login with password <code>forced-commands-only</code> Superuser can only run commands in SSH command line
<code>AllowUsers jdoe ksmith</code> <code>DenyUsers jhacker</code>	List of users that can/cannot login via SSH, or <code>*</code> for everybody
<code>AllowGroups geeks</code> <code>DenyGroups *</code>	List of groups whose members can/cannot login via SSH, or <code>*</code> for all groups
<code>PasswordAuthentication yes</code>	Permit authentication via login and password
<code>PubKeyAuthentication yes</code>	Permit authentication via public key
<code>HostbasedAuthentication yes</code>	Permit authentication based on trusted hosts
<code>Protocol 1,2</code>	Specify protocols supported by SSH. Value can be 1 or 2 or both
<code>X11Forwarding yes</code>	Allow X11 Forwarding

<code>/etc/ssh/ssh_config</code> and <code>~/.ssh/config</code>	SSH client configuration file
<code>Host *</code>	List of hosts to which the following directives will apply, or <code>*</code> for all hosts
<code>StrictHostKeyChecking yes</code>	Ask before adding new host keys to the <code>~/.ssh/known_hosts</code> file, and refuse to connect if the key for a known host has changed. This prevents MITM attacks
<code>GSSAPIAuthentication yes</code>	Support authentication using GSSAPI
<code>ForwardX11Trusted yes</code>	Allow remote X11 clients to fully access the original X11 display
<code>IdentityFile ~/.ssh/id_rsa</code>	User identity file for authentication. Default values are: <code>~/.ssh/identity</code> for protocol version 1 <code>~/.ssh/id_rsa</code> and <code>~/.ssh/id_dsa</code> for protocol version 2

The **X.509** standard defines the format of public key certificates and other related files; it includes cryptographic standards and protocols such as SSL/TLS, PKCS #7, PKCS #12, and OCSP. The **Public Key Infrastructure X.509 (PKIX)** is described in RFC 5280.

X.509 file formats	
DER	Binary-encoded certificate
PEM	ASCII-armored Base64-encoded certificate, included between these two lines: -----BEGIN <i>FILE_TYPE</i> ----- -----END <i>FILE_TYPE</i> ----- where <i>FILE_TYPE</i> is one of the X.509 file types (see below)
DER and PEM are also used as file extensions for different types of files (see below).	

X.509 file type extensions	
CRT CER	Certificate or certificate chain
CSR	Certificate Signing Request
KEY	Private key
CRL	Certificate Revocation List
DER	Certificate; DER-encoded
PEM	Certificate (including or not the private key), certificate chain, or Certificate Signing Request; PEM-encoded

Other file type extensions	
P12 PFX	Certificate (including or not the private key), certificate chain, or Certificate Signing Request; bundled in a PKCS #12 archive file format

OpenSSL is an open source cryptographic library containing an implementation of the SSL (Secure Sockets Layer) and TLS (Transport Layer Security) protocols, plus various cryptographic utilities.

<code>openssl</code>	OpenSSL command line tool
<code>genkey</code>	Interactive utility for the generation of SSL certificates and Certificate Signing Requests
<code>certwatch</code>	Program that issues email warnings when an SSL certificate is close to its expiration date
<code>keyrand</code>	Utility that collects random bits from <code>/dev/random</code> and appends them to a file
<code>CA.pl</code>	User-friendly command for common certificate operations
<code>CA.pl -newca</code>	Create a Certification Authority hierarchy
<code>CA.pl -newreq</code>	Generate a Certificate Signing Request
<code>CA.pl -newreq-nodes</code>	Generate a Certificate Signing Request, creating also a key pair (unencrypted, for non-interactive use)
<code>CA.pl -signreq</code>	Sign a Certificate Signing Request
<code>CA.pl -pkcs12 "Cert name"</code>	Generate a PKCS #12 certificate from a Certificate Signing Request
<code>CA.pl -newcert</code>	Generate a self-signed certificate
<code>CA.pl -verify</code>	Verify a certificate against the Certification Authority certificate for "demoCA"

```
openssl x509 -text -in cert.crt -noout
openssl req -text -in cert.csr -noout
openssl req -new -key cert.key -out cert.csr

openssl req -new -keyout cert.key -out cert.csr \
-newkey rsa:2048 -nodes

openssl x509 -req -in cert.csr -CAcreateserial \
-CA ca.crt -CAkey ca.key -out cert.crt -days validity

openssl req -x509 -keyout cert.key -out cert.crt \
-newkey rsa:2048 -nodes -days validity

openssl ca -config ca.conf -in cert.csr \
-out cert.crt -days validity -verbose

openssl ca -config ca.conf -gencrl -revoke cert.crt \
-crl_reason why

openssl ca -config ca.conf -gencrl -out list.crl
```

```
openssl x509 -in cert.pem -outform DER -out cert.der
openssl pkcs12 -export -in cert.pem \
-inkey cert.key -out cert.pfx -name friendlyname

openssl pkcs12 -in cert.p12 -out cert.crt -clcerts \
-nokeys

openssl pkcs12 -in cert.p12 -out cert.key -nocerts \
-nodes

openssl pkcs12 -in cert.p12 -out ca.crt -cacerts
cat cert.crt cert.key > cert.pem
```

```
openssl dgst -hashfunction -out file.hash file
openssl dgst -hashfunction file | cmp -b file.hash
openssl dgst -hashfunction -sign private.key \
-out file.sig file

openssl dgst -hashfunction -verify public.key \
-signature file.sig file

openssl enc -e -cipher -in file -out file.enc -salt
openssl enc -d -cipher -in file.enc -out file
```

```
openssl genpkey -algorithm RSA -cipher 3des \
-pkeyopt rsa_keygen_bits:2048 -out keypair.pem

openssl pkey -text -in private.key -noout
openssl pkey -in old.key -out new.key -cipher
openssl pkey -in old.key -out new.key
```

1. `openssl s_client -connect www.site.com:443 > tmpfile`
2. **CTRL C**
3. `openssl x509 -in tmpfile -text`

```
openssl list-message-digest-commands
openssl list-cipher-commands
```

Read a certificate

Read a Certificate Signing Request

Generate a Certificate Signing Request, given a private key

Generate a Certificate Signing Request, creating also a 2048-bit RSA key pair (unencrypted, for non-interactive use)

Sign a certificate as a CA, given a Certificate Signing Request

Generate a self-signed root certificate, and create a new CA private key

Sign a certificate

Revoke a certificate

Generate a Certificate Revocation List containing all revoked certificates so far

Convert a certificate from PEM to DER

Convert a certificate from PEM to PKCS #12 including the private key

Convert a certificate from PKCS #12 to PEM

Extract the private key from a PKCS #12 certificate

Extract the CA certificate from a PKCS #12 certificate

Create a PEM certificate from CRT and private key

Generate the digest (hash) of a file

Check the hash of a file; no output means OK

Sign a file

Verify the signature of a file

Encrypt a file

Decrypt a file

Generate a 2048-bit RSA key pair protected by a TripleDES-encrypted passphrase

Examine a private key

Change the passphrase of a private key

Remove the passphrase from a private key

Inspect an SSL certificate from a website

List all available hash functions

List all available ciphers

In **symmetric cryptography**, a symmetric cipher (i.e. cryptographic algorithm) is used with a **shared secret key** to encrypt a message. The message can then be decrypted using the same key.

In **asymmetric cryptography** aka **Public Key Cryptography**, ciphers do not operate with a single key but with a **key pair**, composed of a **public key** and a **private key**. Public and private key are created together at the same time using a special algorithm and are strictly related to each other; however, deriving a private key from its public key is computationally infeasible. A message is encrypted with a public key and can only be decrypted with the companion private key. Similarly, a message is digitally signed with a private key and can only be verified with the companion public key.

Encryption guarantees confidentiality (only sender and receiver are able to know the contents of the message).

Digital signature guarantees authentication (the receiver can verify that the message originates from the intended sender), integrity (the receiver can verify that the message was not modified since it was signed), and non-repudiation (the sender cannot deny having signed the message).

In a **block cipher**, the input plaintext is split into blocks of fixed size, fed to the algorithm, and transformed according to the key. If the input plaintext is shorter than the block length, padding is used.

In a **stream cipher**, the input plaintext is combined with a pseudorandom digit stream. The key is applied to each bit, one at a time.

A **hash function** maps a bitstring of arbitrary length to another bitstring of fixed length, hence outputting a condensed representative image of the bitstring fed in input. Changing just one bit of the input string results in a very different hash value in output (avalanche effect).

A hash function must have the following properties:

- be one-way, i.e. given an output value it is computationally infeasible to find the matching input (preimage resistance)
- given a specific input, it is computationally infeasible to find another input that results in the same output (second preimage resistance)
- it is computationally infeasible to find two different inputs which would result in the same output (collision resistance)

A **Public Key Infrastructure (PKI)** handles creation, management, distribution, use, and revocation of Digital Certificates. It is composed of the following entities:

- the **Certification Authority (CA)** which issues and verifies Digital Certificates
- the **Registration Authority (RA)** which verifies the CA, verifies the subject, and ensures valid and correct registration
- the **Validation Authority (VA)** which verifies the validity of a Digital Certificate

bcrypt	File encryption tool. Uses the Blowfish algorithm
ccrypt	File and stream encryption tool. Uses the Rijndael block cipher
ccr	Codecrypt, an encryption and signing tool that uses only algorithms resistant to quantum-computer cryptanalysis
age	File encryption tool
stegsnow	Steganography tool for text files. The secret message is concealed in additional tab and whitespace characters at the end of lines
steghide	Steganography tool for image and audio files
shasum shasum sha224sum sha256sum sha384sum sha512sum	Print or check the digest of a file generated by different SHA hashing algorithms
md5sum	Print or check the digest of a file generated by the MD5 hashing algorithm
md5pass	Create an MD5 password hash. If no salt is specified, a random salt will be generated

Symmetric ciphers	
DES (Data Encryption Standard)	Block cipher with a 64-bit block size. Uses DEA (Data Encryption Algorithm) with a 56-bit key. Obsolete and insecure.
3DES (Triple DES)	Cipher which uses a key bundle of three DES keys: K_1 , K_2 , and K_3 . The algorithm consists in applying DES three times: encrypt with K_1 , decrypt with K_2 , then encrypt with K_3 . Key options are: - K_1 , K_2 , and K_3 are all different (most secure) - $K_1 = K_3$ - $K_1 = K_2 = K_3$ (fallback to DES, insecure)
AES (Advanced Encryption Standard)	Iterated block cipher with a 128-bit block size. NIST standard. Can use a 128-bit, 192-bit, or 256-bit key.
IDEA (International Data Encryption Algorithm)	Block cipher with a 64-bit block size and a 128-bit key. Intended as a replacement for DES, is now patent-free.
RC4 (Rivest Cipher 4) aka ARC4 or ARCFOUR	Stream cipher with byte-oriented operations, based on the use of a random permutation. Key size is variable. Patented by RSA Security.
RC5 (Rivest Cipher 5)	Fast, parameterized block cipher with a 32-bit, 64-bit, or 128-bit block size. Number of rounds can be up to 255, and key size can be up to 2040 bits. Uses two 2-bit registers. Patented by RSA Security.
RC6 (Rivest Cipher 6)	Block cipher derived from RC5, with the improvement of integer multiplication and four 4-bit registers. Patented by RSA Security.
SEAL (Software-optimized Encryption Algorithm)	Stream cipher optimized for 32-bit machines. Patented by IBM.
Blowfish	16-round Feistel block cipher with a 64-bit block size. Key size can be from 32 to 448 bits, and is expanded into a set of subkeys 4168 bytes long. Designed by Bruce Schneier as a replacement for DES or IDEA.
Twofish	Feistel block cipher with a 128-bit block size. Highly configurable. Designed by Bruce Schneier et al.
Asymmetric ciphers	
DSA (Digital Signature Algorithm)	NIST standard used in the Digital Signature Standard (DSS). Generates a 320-bit digital signature. Based on the algebraic properties of modular exponentiation and the discrete logarithm problem.
RSA (Rivest-Shamir-Adleman)	Cryptosystem for encryption and authentication. Based on modular arithmetic and large prime numbers.
DH (Diffie-Hellman)	Cryptographic protocol for securely establishing a shared secret key over an insecure channel. Based on modular arithmetic, exponentiation, and large prime numbers. Strength and algorithm of key exchange are defined by groups: - group 1 (768 bit) - group 2 (1024 bit) - group 5 (1536 bit) - group 14 (2048 bit) - group 19 (256 bit, elliptic curve) - group 20 (384 bit, elliptic curve)

Hash functions	
MD2 (Message Digest v2)	Takes in input a message which is a multiple of 512 bits (if not, padding is used) and has a maximum length of $2^{64}-1$ bits, and generates a 128-bit hash. Supports 8-bit machines (i.e. word size is 8 bits). Broken, but still used in digital certificates together with RSA.
MD4 (Message Digest v4)	Same properties as MD2. Supports 32-bit machines. Proven severely insecure, thus obsolete.
MD5 (Message Digest v5)	Same properties as MD2. Supports 32-bit machines. This is the hash function of the MD family currently in use.
SHA-0 (Secure Hash Algorithm v0)	Original version of the SHA hash function. Vulnerable, thus not in use anymore.
SHA-1 (Secure Hash Algorithm v1)	Takes in input a message with a maximum length of $2^{64}-1$ bits, and generates a 160-bit hash. Vulnerable and therefore no longer approved for cryptographic use.
SHA-2 (Secure Hash Algorithm v2)	Family of hash functions: SHA-256 (for 32-bit machines, generates a 256-bit hash) SHA-224 (for 32-bit machines, generates a 224-bit hash, truncated version) SHA-512 (for 64-bit machines, generates a 512-bit hash) SHA-384 (for 64-bit machines, generates a 384-bit hash, truncated version) Max input message length is $2^{128}-1$ bits.
SHA-3 (Secure Hash Algorithm v3)	Family of hash functions: SHA3-256 (for 32-bit machines, generates a 256-bit hash) SHA3-224 (for 32-bit machines, generates a 224-bit hash, truncated version) SHA3-512 (for 64-bit machines, generates a 512-bit hash) SHA3-384 (for 64-bit machines, generates a 384-bit hash, truncated version) SHAKE128 (generates a hash of arbitrary length) SHAKE256 (generates a hash of arbitrary length) Max input message length is unlimited.
RIPEMD (RACE Integrity Primitives Evaluation Message Digest)	Family of hash functions: RIPEMD (generates a 128-bit hash; insecure) RIPEMD-128 (generates a 128-bit hash; insecure) RIPEMD-160 (generates a 160-bit hash; most commonly used) RIPEMD-256 (generates a 256-bit hash) RIPEMD-320 (generates a 320-bit hash)

Authentication systems	
HMAC (Hash-based Message Authentication Code)	Message Authentication Code used to verify data integrity and sender authentication. Uses a hash function in conjunction with a secret key.
PAP (Password Authentication Protocol)	Password-based authentication protocol used by Point-to-Point Protocol (PPP) to authenticate remote clients. Uses a weak authentication scheme, vulnerable to attacks; passwords are transmitted in plaintext over the network. For this reason, is not recommended and it is advised to use CHAP or EAP instead.
CHAP (Challenge-Handshake Authentication Protocol)	Authentication protocol used by Point-to-Point Protocol (PPP) to authenticate remote clients. Client identity is verified via a three-way handshake. It uses an incrementally changing identifier and a variable challenge value in order to thwart replay attacks.
EAP (Extensible Authentication Protocol)	Authentication framework able to use different authentication systems (passwords, smart tokens, one-time passwords, Secure ID cards, digital certificates, public key cryptography protocols, etc.) via a challenge-response mechanism.
LEAP (Lightweight Extensible Authentication Protocol)	Cisco proprietary version of EAP, used for WEP. Uses either the MS-CHAP or the EAP-FAST authentication protocol. Vulnerable and not recommended.
PEAP (Protected Extensible Authentication Protocol)	TLS-encapsulated secured version of EAP, used in WPA2.

PGP (Pretty Good Privacy), created in 1991 by Phil Zimmermann, is the first tool for strong encryption available to the public. From it derived the **OpenPGP standard**, described in RFC 4880 and which specifies a suite of algorithms: ElGamal, DSA, Triple DES, SHA-1, RSA, AES-128, CAST-128, IDEA, Camellia, ECC (Elliptic Curve Cryptography) i.e. ECDSA and ECDH.

GPG (GNU Privacy Guard) aka **GnuPG** is a well-known implementation of OpenPGP.

<code>gpg --gen-key</code>	Generate a key pair
<code>gpg --import <i>alice.asc</i></code>	Import Alice's public key <i>alice.asc</i> into your keyring
<code>gpg --list-keys</code>	List the keys contained into your keyring
<code>gpg --list-secret-keys</code>	List your private keys contained into your keyring
<code>gpg --list-public-keys</code>	List the public keys contained into your keyring
<code>gpg --export -o <i>keyring.gpg</i></code>	Export your whole keyring to a file <i>keyring.gpg</i>
<code>gpg --export-secret-key -a "You" -o <i>private.key</i></code>	Export your private key to a file <i>private.key</i>
<code>gpg --export-public-key -a "Alice" -o <i>alice.pub</i></code>	Export Alice's public key to a file <i>alice.pub</i>
<code>gpg --edit-key "Alice"</code>	Sign Alice's public key
<code>gpg -e -u "You" -r "Alice" <i>file</i></code>	Sign <i>file</i> (with your private key) and encrypt it to Alice (with Alice's public key)
<code>gpg -d <i>file.gpg</i> -o <i>file</i></code>	Decrypt <i>file.gpg</i> (with your own private key) and save the decrypted file to <i>file</i>

<code>pinentry</code>	Helper dialog programs that allow GPG to securely read passphrases and PINs
-----------------------	---

<code>dirmngr</code>	Server for managing CRLs and certificates
----------------------	---

<code>dirmngr-client</code>	Client for <code>dirmngr</code>
-----------------------------	---------------------------------

<code>gpgtar</code>	Archive files, signing or encrypting them at the same time
---------------------	--

<code>gpgv</code>	Verify a signature
-------------------	--------------------

<code>watchgnupg</code>	Socket listener to read GPG logs
-------------------------	----------------------------------

<code>addgnupghome</code>	Create GPG home directories
---------------------------	-----------------------------

<code>gpgconf</code>	Modify GPG configuration in home directories
----------------------	--

<code>applygnupgdefaults</code>	Modify GPG configuration in home directories, enforcing policies for all users
---------------------------------	--

<code>gpg-check-pattern</code>	Check a passphrase against a pattern file
--------------------------------	---

<code>gpg-preset-passphrase</code>	Seed the <code>gpg-agent</code> cache with passphrases
------------------------------------	--

<code>gpg-connect-agent</code>	Connect with a <code>gpg-agent</code>
--------------------------------	---------------------------------------

Sequoia is an alternative implementation of OpenPGP.

<code>sq</code>	Command line frontend for Sequoia
-----------------	-----------------------------------

<code>sqv</code>	OpenPGP signature verification tool
------------------	-------------------------------------

<code>sq-git</code>	Tool to protect a Git repository
---------------------	----------------------------------

LUKS (Linux Unified Key Setup) is a platform-independent specification for the encryption of a block device. It employs **dm-crypt**, a transparent disk encryption subsystem which is part of the device mapper, as backend. The `lsblk` command, that lists devices and partitions, can be used to identify LUKS-encrypted ones.

<code>cryptsetup</code>	Frontend command for dm-crypt. Will prompt for a passphrase for most operations on a LUKS-encrypted device
<code>cryptsetup luksFormat device</code>	Initialize a LUKS partition, prompting for an encryption passphrase
<code>cryptsetup luksChangeKey device</code>	Change the passphrase of a LUKS partition
<code>cryptsetup luksAddKey device</code>	Add a new passphrase to a LUKS partition
<code>cryptsetup luksAddKey device keyfile</code>	Add a new keyfile to a LUKS partition
<code>cryptsetup luksRemoveKey device</code>	Remove a passphrase from a LUKS partition
<code>cryptsetup luksRemoveKey device keyfile</code>	Remove a keyfile from a LUKS partition
<code>cryptsetup luksKillSlot device keyslot</code>	Remove a key from a LUKS partition
<code>cryptsetup isLuks device</code>	Return true if the device is a LUKS partition
<code>cryptsetup luksDump device</code>	Dump the header information of a LUKS partition
<code>cryptsetup luksUUID device</code>	Print the UUID a LUKS partition
<code>cryptsetup luksOpen device name</code>	Open a LUKS device and set up a mapping name
<code>cryptsetup luksClose name</code>	Close a LUKS device and remove the mapping name
<code>cryptsetup luksSuspend name</code>	Suspend a LUKS device and wipe the encryption key from memory
<code>cryptsetup luksResume name</code>	Resume a suspended LUKS device
<code>cryptsetup luksHeaderBackup device --header-backup-file file</code>	Backup header and keyslot areas of a LUKS device to a file
<code>cryptsetup luksHeaderRestore device --header-backup-file file</code>	Restore header and keyslot areas of a LUKS device from a file

OpenVPN is an open source software that implements a Virtual Private Network (VPN) between two endpoints. The encrypted VPN tunnel uses UDP port 1194.

`openvpn --genkey --secret keyfile` Generate a shared secret keyfile for OpenVPN authentication.
The keyfile must be copied on both server and client

`openvpn server.conf` Start the VPN on the server side

`openvpn client.conf` Start the VPN on the client side

`/etc/openvpn/server.conf` Server-side configuration file:

```
dev tun
ifconfig server_IP client_IP
keepalive 10 60
ping-timer-rem
persist-tun
persist-key
secret keyfile
```

`/etc/openvpn/client.conf` Client-side configuration file:

```
remote server_public_IP
dev tun
ifconfig client_IP server_IP
keepalive 10 60
ping-timer-rem
persist-tun
persist-key
secret keyfile
```

<code>last</code>	Print the list of users that logged in and out. Searches through the file <code>/var/log/wtmp</code>
<code>last -F user</code>	Print last login times of <i>user</i> with full date (including the year)
<code>lastb</code>	Print the list of bad login attempts. Searches through the file <code>/var/log/btmp</code>
<code>faillog -a</code>	Display failed login attempts from all users. Searches through the file <code>/var/log/faillog</code>
<code>faillog -u user</code>	Display failed login attempts from <i>user</i>
<code>faillog -l 60 -m 5</code>	Lock an account for 60 seconds after a failed login, and disable it after 5 failed logins
<code>pam_tally2</code>	Deny access to users that have too many failed logins
<code>acct on</code> <code>acct off</code>	Turn process accounting on or off
<code>ac</code>	Print statistics about connect time of users
<code>lastcomm</code>	Print information about previously executed commands
<code>sa</code>	Print summarized information about previously executed commands

Fail2Ban is an intrusion prevention software framework. It temporarily bans IP addresses that e.g. have too many failed password logins (by analyzing authentication logs) or cause too many HTTP errors (by analyzing webserver's logs).

<code>fail2ban-client</code>	Manage the Fail2Ban system
<code>fail2ban-client status</code>	Show a summarized status of all jails
<code>fail2ban-client status jail</code>	Show the status of <i>jail</i>
<code>fail2ban-client status --all</code>	Show the status of all jails
<code>fail2ban-client stats</code>	Display statistics about all jails

`/etc/fail2ban/filter.d/myjail1-error404.conf`

```
[Definition]
failregex = ^<HOST>.*"GET.*" 404 .*$
ignoreregex =
```

`/etc/fail2ban/filter.d/myjail2-scriptunknown.conf`

```
[Definition]
failregex = ^.* \[client <HOST>:\d+\] AH01071: Got error \'Primary script unknown\''.*$
ignoreregex =
```

`/etc/fail2ban/jail.local`

```
[myjail1-error404]
enabled = true
filter = myjail1-error404
logpath = /var/log/httpd/access_log
maxretry = 7
findtime = 30
bantime = 86400
action = firewallcmd-rich-rules[actiontype=<multiport>, port="80,443", \
    protocol="tcp", rich-blocktype="reject type='icmp-port-unreachable'"]
```

This jail applies a filter to the webserver access log to find the IP addresses that cause an HTTP 404 error. It then bans HTTP and HTTPS access for 24 hours to the IP addresses hitting 7 errors within 30 seconds. The ban is applied via the creation of a firewall rich rule.

```
[myjail2-scriptunknown]
enabled = true
filter = myjail2-scriptunknown
logpath = /var/log/httpd/error_log
maxretry = 5
findtime = 60
bantime = 43200
action = firewallcmd-rich-rules[actiontype=<multiport>, port="80,443", \
    protocol="tcp", rich-blocktype="reject type='icmp-port-unreachable'"]
```

This jail applies a filter to the webserver error log to find the IP addresses that cause an error of type AH01071, which identifies an attempt to run a PHP/FastCGI script. It then bans HTTP and HTTPS access for 12 hours to the IP addresses hitting 5 errors within 60 seconds. The ban is applied via the creation of a firewall rich rule.

Both jails have the purpose of blocking clients that perform vulnerability scans on the webserver.

Key	Alternate key	Function
CTRL F	→	Move cursor forward one character
CTRL B	←	Move cursor backward one character
CTRL A	HOME	Move cursor to beginning of line
CTRL E	END	Move cursor to end of line
CTRL H	BACKSPACE	Delete character to the left of cursor
CTRL W		Delete word to the left of cursor
CTRL U		Delete all characters to the left of cursor
CTRL K		Delete all characters to the right of cursor
CTRL T		Swap current character with previous one
ESC T		Swap current word with previous one
SHIFT PAGE UP		Scroll up the screen buffer
SHIFT PAGE DOWN		Scroll down the screen buffer
CTRL L		Clear screen (same as <code>clear</code>)
CTRL P	↑	Previous command in history
CTRL N	↓	Next command in history
CTRL R		Reverse history search
ALT .		Get argument from previous command in history
CTRL I	TAB	Autocomplete commands, filenames, and directory names
ALT /		Autocomplete filenames and directory names only
CTRL ALT E		Expand the Bash alias currently entered on the command line
CTRL J	RETURN	Line feed
CTRL M		Carriage return
CTRL S		Pause transfer to terminal Forward history search (if XON/XOFF flow control is disabled)
CTRL Q		Resume transfer to terminal
CTRL Z		Send a SIGTSTP to put the current job in background
CTRL C		Send a SIGINT to stop the current process
CTRL D		Send an EOF to the current process (same as <code>logout</code> if process is a shell)
CTRL ALT DEL		Send a SIGINT to reboot the machine (same as <code>shutdown -r now</code>), as specified in <code>/etc/inittab</code> and <code>/etc/init/control-alt-delete</code>
CTRL ALT F1 ... F6		Switch between text consoles (same as <code>chvt n</code>)

Key	Alternate key	Function
CTRL ALT F7 ... F11		Switch between X Window consoles
CTRL ALT +		Increase X Window screen resolution
CTRL ALT -		Decrease X Window screen resolution
CTRL TAB		Switch between X Window tasks
CTRL ALT →	CTRL ALT ↓	Switch to next workspace
CTRL ALT ←	CTRL ALT ↑	Switch to previous workspace
CTRL ALT BACKSPACE		Reboot the X Window server
GNOME		
ALT TAB		Switch between windows in the current workspace
SUPER		Show activities overview
SUPER L		Lock screen
SUPER M		Show tray messages
SUPER ↑		Maximize current window
SUPER ↓		Restore normal size of current window
SUPER ←		Maximize current window to left half screen
SUPER →		Maximize current window to right half screen
ALT F2		Run command
CTRL +		Increase terminal font size
CTRL -		Decrease terminal font size

The **kernel** is a computer program at the core of the Linux (or any other) OS. It is always memory-resident, runs at the lowest level, manages processes, handles interactions between hardware and software, and has always full control over the operating system.

A kernel version number has the form *kernelversion.majorrevision.minorrevision*.
Kernel versions of the form *X.Y.Z-A.B.C.foo* are kernels packaged and modified for a specific distribution.
Kernel images are usually gzip-compressed and can be of two types: zImage (max 520 Kb) and bzImage (no size limit).

Linux and UNIX are OSes with a **monolithic kernel** architecture. In this design, the entire system runs in kernel space. The kernel consists of a single, large executable file whose functions are memory management, network management, device drivers, file system management, and process management. The monolithic kernel bypasses the need for IPC mechanisms, thus simplifying communication and reducing overhead.
The **microkernel** architecture (e.g MINIX), on the other hand, consists of a kernel with only core functionalities such as process and memory management; other services such as device drivers, file system management, and networking, run as separate processes in user space. This ensures modularity and extensibility.

/boot/vmlinuz-5.14.0-427.13.1.el9_4.x86_64	Linux compressed kernel file: version 5.14.0, patch level 427.13.1, RHEL 9.4 (or a derivative), 64-bit x86 architecture. These values are the same outputted by <code>uname -r</code>
/usr/src/linux/	Directory containing the kernel source code to be compiled
/usr/src/linux/.config	Kernel configuration file
/etc/initramfs-tools/initramfs.conf (Debian)	Configuration file for the <code>initrd</code> image file

The runtime loader `ld.so` loads the required shared libraries of the program into RAM, searching in this order:

- 1. `LD_LIBRARY_PATH` Environment variable specifying the list of dirs where libraries should be searched for first
- 2. `/etc/ld.so.cache` Cache file
- 3. `/lib` and `/usr/lib` Default locations for shared libraries
- 4. Other shared library locations specified in the file `/etc/ld.so.conf`

Configuration of device drivers	
Device drivers support the kernel with instructions on how to use that device.	
Device driver compiled into the kernel	Configure the device driver by passing a kernel parameter in the GRUB menu: kernel /vmlinuz ro root=/dev/vg0/root vga=0x33c
Device driver provided as a kernel module	Edit module configuration in <code>/etc/modprobe.conf</code> or <code>/etc/modprobe.d/</code> (Red Hat): alias eth0 3c59x Specify that eth0 uses the 3c59x.ko driver module options 3c509 irq=10,11 Assign IRQ 10 and 11 to 3c509 devices

<code>freeramdisk</code>	Free the memory used for the <code>initrd</code> image. This command must be run directly after unmounting <code>/initrd</code>
<code>mkinitrd <i>initrd_image</i> <i>kernel_version</i></code> (Red Hat) <code>mkinitramfs</code> (Debian)	Create an <code>initrd</code> image file
<code>dracut</code>	Create initial ramdisk images for preloading modules
<code>lsinitramfs</code>	Show the contents of an <code>initramfs</code> image
<code>dbus-monitor</code>	Monitor messages going through a D-Bus message bus
<code>dbus-monitor --session</code>	Monitor session messages (default)
<code>dbus-monitor --system</code>	Monitor system messages
<code>kexec -l <i>kernel_image</i> --append=<i>options</i> \</code> <code>--initrd=<i>initrd_image</i> && kexec -e</code>	Load a kernel image file into memory and boot it. This allows running a different kernel without rebooting the machine
<code>ldconfig</code>	Create a cache file <code>/etc/ld.so.cache</code> of all available dynamically linked libraries. This command should be run when the system complains about missing libraries
<code>ldd <i>program_or_lib</i></code>	Print library dependencies

<code>lspci</code>	List PCI devices
<code>lspci -d 8086:</code>	List all Intel hardware present. PCI IDs are stored in: /usr/share/hwdata/pci.ids (Red Hat) /usr/share/misc/pci.ids (Debian)
<code>lsusb</code>	List USB devices
<code>lsusb -d 8086:</code>	List all Intel USB devices present. USB IDs are stored in: /usr/share/hwdata/usb.ids (Red Hat) /var/lib/usbutils/usb.ids (Debian)
<code>lspcmcia</code>	List PCMCIA devices
<code>lsdev</code>	List information about the system hardware
<code>lshw</code>	List system hardware
<code>lscpu</code>	List information about the CPU architecture
<code>dmidecode</code>	List information about the system's hardware, serial number, BIOS version, and other data. This tool produces a dump of the DMI (Desktop Management Interface) table, aka SMBIOS (System Management BIOS) table, in human-readable format
<code>uname</code>	Print system information
<code>uname -s</code>	Print the kernel name
<code>uname -n</code>	Print the network node hostname
<code>uname -r</code>	Print the full kernel release number, distro, and architecture
<code>uname -v</code>	Print the kernel version number
<code>uname -m</code>	Print the machine hardware name
<code>uname -p</code>	Print the processor type
<code>uname -i</code>	Print the hardware platform
<code>uname -o</code>	Print the operating system
<code>uname -a</code>	Print all the above information, in that order
<code>cpufetch</code>	List information about the CPU architecture as well as an ASCII art image of the vendor
<code>evtest</code>	Monitor and query input device events in <code>/dev/input/eventn</code>
<code>stap script.stp</code>	Run a SystemTap script. SystemTap is a RHEL tool to gather information on the kernel
<code>stap-prep</code>	Set up the SystemTap environment
<code>stap-report</code>	Verify the SystemTap environment

Kernel compile		
Download	Download the kernel source code <code>linux-X.Y.Z.tar.bz2</code> from http://www.kernel.org to the base of the kernel source tree <code>/usr/src/linux</code>	
Clean	<code>make clean</code>	Delete most generated files
	<code>make mrproper</code>	Delete all generated files and kernel configuration
	<code>make distclean</code>	Delete temporary files, patch leftovers, and similar files
Configure	<code>make config</code>	Create configuration (terminal-based; options must be set in sequence)
	<code>make menuconfig</code>	Create configuration (ncurses UI)
	<code>make xconfig</code> <code>make gconfig</code>	Create configuration (GUI)
	<code>make oldconfig</code>	Create a new configuration file, based on the options in the old configuration file and in the source code
	Components (e.g. device drivers) can be either: - not compiled - compiled into the kernel binary, for support of devices always used on the system or necessary for the system to boot - compiled as a kernel module, for optional devices The configuration command creates a configuration file <code>/usr/src/linux/.config</code> containing instructions for the kernel compilation	
Build	<code>make bzImage</code>	Compile the kernel
	<code>make modules</code>	Compile the kernel modules
	<code>make all</code>	Compile kernel and kernel modules
	<code>make -j2 all</code> will speed up compilation by allocating 2 simultaneous compile jobs	
Modules install	<code>make modules_install</code>	Install the previously built modules present in <code>/lib/modules/X.Y.Z</code>
Kernel install	<code>make install</code>	Install the kernel automatically
	To install the kernel by hand: - Copy the new compiled kernel and other files into the boot partition: <code>cp /usr/src/linux/arch/boot/bzImage /boot/vmlinuz-X.Y.Z</code> (kernel) <code>cp /usr/src/linux/arch/boot/System.map-X.Y.Z /boot</code> <code>cp /usr/src/linux/arch/boot/config-X.Y.Z /boot</code> (config options used for this compile) - Create an entry in GRUB to boot on the new kernel	
Package	Optionally, the kernel can be packaged for install on other machines	
	<code>make rpm-pkg</code>	Build source and binary RPM packages
	<code>make binrpm-pkg</code>	Build binary RPM package
	<code>make deb-pkg</code>	Builds binary DEB package

Kernel patching		
Download	Download and decompress the patch to <code>/usr/src</code>	
Patch	<code>patch -p1 < file.patch</code>	Apply the patch
	<code>patch -Rp1 < file.patch</code>	Remove (reverse) a patch. Alternatively, applying the patch again reverses it
Build	Build the patched kernel as explained above	
Install	Install the patched kernel as explained above	

The monolithic kernel architecture solves the disadvantage of lack of modularity and extensibility by using kernel **modules**, which are object files that contain executable code and allow the kernel to access functions (symbols) for kernel services e.g. hardware drivers, network stack, or filesystem abstraction.

Kernel modules can be loaded dynamically into the kernel to provide additional functionalities on demand at runtime, instead of being included when the kernel is compiled; this reduces memory footprint. They can then be unloaded when no longer necessary.

`kerneld` (daemon) and `kmod` (kernel thread) facilitate the dynamic loading of kernel modules.

<code>/lib/modules/X.Y.Z/*.ko</code>	Kernel modules for kernel version <code>X.Y.Z</code>
<code>/lib/modules/X.Y.Z/modules.dep</code>	Modules dependencies. This file needs to be recreated (via <code>depmod -a</code>) after a reboot or a change in module dependencies
<code>/etc/modules.conf</code> <code>/etc/conf.modules</code> (deprecated)	Modules configuration file
<code>lsmod</code>	List the modules that are currently loaded into the kernel
<code>insmod module</code>	Insert a module into the kernel. Will fail if the module requires another module or if compatible hardware is not detected. It is better to use <code>modprobe</code> instead
<code>rmmod module</code>	Remove a module from the kernel. If the module is in use by another module, the latter module must be removed first. It is better to use <code>modprobe -r</code> instead
<code>modinfo module</code>	Display the list of parameters accepted by the module
<code>depmod -a</code>	Probe all modules in the kernel modules directory and generate the file that lists their dependencies
<code>modprobe module option=value</code>	Insert a module into the running kernel, with the specified parameters. Prerequisite modules will be inserted automatically. It is recommended to use <code>modprobe</code> instead of <code>insmod</code> and <code>rmmod</code> , because it automatically handles prerequisites when inserting modules, is more specific about errors, and accepts just the module name alone instead of requiring the full path
<code>modprobe -a</code>	Insert all modules
<code>modprobe -t directory</code>	Attempt to load all modules contained in the directory until a module succeeds. This action probes the hardware by successive module-insertion attempts for a single type of hardware, e.g. a network adapter
<code>modprobe -r module</code>	Remove a module
<code>modprobe -c module</code>	Display module configuration
<code>modprobe -l</code>	List loaded modules
<code>find /lib/modules/\$(uname -r) -type f -name '*.ko*'</code>	List all loadable modules (including compressed ones)

<code>dmesg</code>	Print the messages of the kernel ring buffer. Each entry is prepended by a timestamp showing the number of seconds since the machine booted up
<code>dmesg -T</code>	Print kernel messages with human-readable timestamps
<code>dmesg -n 1</code>	Print kernel messages with logging level 1 (i.e. only panic messages)
<code>journalctl</code>	Display the Systemd journal, which contains the kernel logs
<code>journalctl -n n</code>	Display the most recent <i>n</i> log lines (default is 10)
<code>journalctl --since "1 hour ago"</code>	Display only events happened in the last hour
<code>journalctl --since "2026-01-13 23:58:00" \</code> <code>--until "2026-01-17 05:00:00"</code>	Display only events happened between two datetimes
<code>journalctl -x</code>	Display events, adding explanations from the message catalog
<code>journalctl -f</code>	Display the journal in real-time
<code>journalctl -u crond.service</code> <code>journalctl _SYSTEMD_UNIT=crond.service</code>	Display the log entries created by the cron service
<code>journalctl --rotate && \</code> <code>journalctl --vacuum-time=1s</code>	Remove all current journal entries
<code>mkdir -p /var/log/journal/ && \</code> <code>systemctl restart systemd-journald</code>	Enable persistent storage of logs in <code>/var/log/journal/</code> (by default, <code>journalctl</code> stores the logfiles in RAM only)

/proc is a pseudo filesystem that gives access to process data held in the kernel.

File	Information stored (can be viewed via <code>cat</code>)	Equivalent command
/proc/bus	Buses (e.g. PCI, USB, PC Card)	
/proc/cpuinfo	CPUs information	
/proc/devices	Drivers currently loaded	
/proc/dma	DMA channels in use	
/proc/filesystems	Filesystems supported by the system	
/proc/interrupts	Current IRQs (Interrupt Requests)	<code>procinfo</code>
/proc/ioports	I/O addresses in use	
/proc/kcore	Memory allocatable by the kernel	
/proc/loadavg	System load averages	<code>uptime</code>
/proc/mdstat	Information about RAID arrays and devices	
/proc/meminfo	Total and free memory	<code>free</code>
/proc/modules	Kernel modules currently loaded	<code>lsmod</code>
/proc/mounts	Mounted partitions	<code>mount</code>
/proc/net/dev	Network interface statistics	
/proc/partitions	Drive partition information	<code>fdisk -l</code>
/proc/swaps	Size of total and used swap areas	<code>swapon -s</code>
/proc/sys/ /proc/sys/kernel/ /proc/sys/net/	sysfs: exposes tunable kernel parameters Kernel information and parameters Network information and parameters	
/proc/uptime	Time elapsed since boot	<code>uptime</code>
/proc/version	Linux version	<code>uname -a</code>
/proc/n/ /proc/n/cmdline /proc/n/cwd /proc/n/envIRON /proc/n/exe /proc/n/fd /proc/n/root /proc/n/status	Information about process with PID <i>n</i> Command by which the process was launched Symlink to process' working directory Values of environment variables of process Symlink to process' executable Files currently opened by the process Symlink to process' filesystem root Status of process	<code>ps n</code> <code>ls -l -p n</code>

/proc/sys is the only writable branch of /proc and can be used to tune kernel parameters on the fly. All changes are lost after system shutdown, unless applied via `sysctl -p`.

```
sysctl fs.file-max
cat /proc/sys/fs/file-max
```

Get the maximum allowed number of open files

```
sysctl -w "fs.file-max=100000"
echo "100000" > /proc/sys/fs/file-max
```

Set the maximum allowed number of open files to 100000

```
sysctl -a
```

List all available kernel tuning options

```
sysctl -p
```

Apply all tuning settings listed in `/etc/sysctl.conf`. This command is usually run at boot by the system initialization script, to make permanent changes to kernel parameters

/dev contains **device files** aka **special files** that represent all devices attached to the local system.

Hardware equipment connected to the machine is represented by a specific **hardware device** file, which allows access to it.

Device file	Hardware device
/dev/sda	SCSI, PATA, or SATA hard drive
/dev/hda	IDE hard drive
/dev/pda	Parallel port IDE hard drive
/dev/vda	Virtual disk for KVM-based virtual machines
/dev/sda, /dev/sdb, /dev/sdc ...	First, second, third ... hard drive
/dev/sda1, /dev/sda2, /dev/sda3 ...	First, second, third ... partition of the first hard drive
/dev/md0	Metadisk group, for use with RAID
/dev/sr0	SCSI CD-ROM
/dev/pcd0	Parallel port CD-ROM
/dev/cdrom	CD-ROM. Usually symlinked to /dev/sr0
/dev/fd0	Floppy disk drive
/dev/ht0	IDE tape drive
/dev/pt0	Parallel port tape drive
/dev/sg0	Generic SCSI device
/dev/lp0	Parallel port printer device
/dev/parport0	Raw parallel port device
/dev/video0	Webcam or other video acquisition device
/dev/autofs	AutoFS device
/dev/fuse	FUSE device
/dev/dsp	Digital Signal Processor device. Interfaces with the soundcard
/dev/fb0	Framebuffer device. Interfaces with the graphics hardware

Device files that correspond to hardware devices can be character special files or block special files.

A driver communicates with a **character device** (terminal, serial port, sound card, mouse, etc.) by single characters (bytes) and via unbuffered access.

A driver communicates with a **block device** (hard disk, USB drive, etc.) by whole blocks of data (usually 4 Kb) and through a cache.

A device file is assigned a unique identifier called **major number**, which is used by the kernel to determine which driver to use when accessing the device.

The kernel can uniquely identify a specific device within a class of devices via its **minor number** in conjunction with the major number. Minor numbers are smaller identifiers used to identify specific devices within a major device class, e.g. different partitions on a hard disk, or different USB drives.

A device file which is not attached to any hardware is a **virtual device**. This allows to emulate a peripheral device.

Device file	Virtual device
/dev/loop0	Loopback pseudo device. Makes a file accessible as a block device, hence allowing a file containing an entire filesystem to be mounted as if it were a disk device
/dev/mem	Physical memory
/dev/kmem	Kernel virtual memory
/dev/core	Obsolete. Symlink to /proc/kcore
/dev/shm	Shared memory device (tmpfs). Can be used like /tmp to store temporary files, but is bound by the amount of RAM in the system
/dev/stdin	Standard Input
/dev/stdout	Standard Output
/dev/stderr	Standard Error
/dev/null	Null device, aka blackhole or bit bucket. Discards any received data
/dev/zero	Zero device. Outputs an infinite stream of zero bytes (NUL) on reads
/dev/full	"Always full" device. Similar to /dev/zero, and also returns an error "No space left on device" (ENOSPC) on writes
/dev/random	Non-deterministic random number generator. Gathers entropy from the system to generate randomness. In old kernels, once the entropy pool was depleted, the device blocked all reads until it could collect more entropy
/dev/urandom	Unlimited pseudo random number generator. Faster but unsafe for cryptographic purposes
/dev/console	System console
/dev/tty	Terminal for current process
/dev/tty0	Current virtual console
/dev/ttyS0	Serial port, usually used for modem connections
/dev/ptyp0	Pseudo-TTY master
/dev/ttyp0	Pseudo-TTY slave

udev is the device manager for the Linux kernel. It dynamically generates the device nodes in `/dev` for devices present on the system; it also provides persistent naming for storage devices in `/dev/disk`.

When a device is added, removed, or changes state, the kernel sends an uevent received by the `udev` daemon which will pass the uevent through a set of rules stored in `/etc/udev/rules.d/*.rules` and `/lib/udev/rules.d/*.rules`.

The **Hardware Abstraction Layer (HAL)** manages device files and provides plug-and-play facilities. The HAL daemon `hald` maintains a persistent database of devices.

<code>udevadm monitor</code> <code>udevmonitor</code>	Show all kernel uevents and udev messages
<code>udevadm info --attribute-walk --name=/dev/sda</code>	Print all attributes of device <code>/dev/sda</code> in udev rules key format
<code>cat /sys/block/sda/size</code>	Print the size attribute of disk <code>sda</code> in 512-byte blocks. This information is retrieved from <code>sysfs</code>
<code>udevadm test /dev/sdb</code>	Simulate an udev event run for the device and print debug output
<code>gnome-device-manager</code>	Browser for the HAL device manager
<code>lshal</code>	Show items in the HAL device database

`/etc/udev/rules.d/*.rules` and `/lib/udev/rules.d/*.rules` **udev rules**

<code>KERNEL=="hda", NAME="mydisk"</code>	Match a device which was named by the kernel as <code>hda</code> ; name the device node as "mydisk". The device node will be therefore <code>/dev/mydisk</code>
<code>KERNEL=="hdb", DRIVER=="ide-disk", SYMLINK+="mydisk myhd"</code>	Match a device with kernel name and driver as specified; name the device node with the default name and create two symbolic links <code>/dev/mydisk</code> and <code>/dev/myhd</code> pointing to <code>/dev/hdb</code>
<code>KERNEL=="fd[0-9]*", NAME="floppy/%n", SYMLINK+="%k"</code>	Match all floppy disk drives (i.e. <code>fdn</code>); place device node in <code>/dev/floppy/n</code> and create a symlink <code>/dev/fdn</code> to it
<code>SUBSYSTEM=="block", ATTR{size}=="41943040", SYMLINK+="mydisk"</code>	Match a block device with a size attribute of 41943040; create a symlink <code>/dev/mydisk</code>
<code>KERNEL=="fd[0-9]*", OWNER="jdoe"</code>	Match all floppy disk drives; give ownership of the device file to user "jdoe"
<code>KERNEL=="sda", PROGRAM="/bin/mydevicenamer %k", SYMLINK+="%c"</code>	Match a device named by the kernel as <code>sda</code> ; to name the device, use the defined program which takes on stdin the kernel name and output on stdout e.g. <code>name1 name2</code> . Create symlinks <code>/dev/name1</code> and <code>/dev/name2</code> pointing to <code>/dev/sda</code>
<code>KERNEL=="sda", ACTION=="add", RUN+="/bin/myprogram"</code>	Match a device named by the kernel as <code>sda</code> ; run the defined program when the device is connected
<code>KERNEL=="sda", ACTION=="remove", RUN+="/bin/myprogram"</code>	Match a device named by the kernel as <code>sda</code> ; run the defined program when the device is disconnected

`%n` = kernel number (e.g. = 3 for `fd3`)

`%k` = kernel name (e.g. = `fd3` for `fd3`)

`%c` = device name as output from program

If the kernel has booted in emergency mode and `init` has not run, some initial configuration is necessary e.g.

```
mount /proc
mount -o remount,rw /
mount -a
```

If mounting the filesystems fails:

```
mknod /dev/sda
mknod /dev/sda1
fdisk -l /dev/sda
fsck -y /dev/sda1
mount -t ext3 /dev/sda1 /mnt/sysimage
chroot /mnt/sysimage
```

To install a package using an alternative root directory (useful if the system has been booted from a removable media):

```
rpm -U --root /mnt/sysimage package.rpm
```

To install GRUB on the specified directory (which must contain `/boot/grub/`):

```
grub-install --root-directory=/mnt/sysimage /dev/sda
```

Alternative method:

```
chroot /mnt/sysimage && grub-install /dev/sda
```

It is a good idea to run `sync` and unmount all filesystems before exiting the shell, to ensure that all changes have been written to disk.

How to reset the root password (RHEL 7 and later)

1. Power up the system and, once on the GRUB 2 boot screen, press **E** to edit the current entry
2. On the kernel line that mentions `linux16`, remove the `rhgb` and `quiet` parameters and add `rd.break` at the end
3. Press **CTRL** **X**; the system will boot on the `initramfs switch_root` prompt
4. Remount the filesystem as writable

```
mount -o remount,rw /sysroot
```
5. Change the filesystem root

```
chroot /sysroot
```
6. Modify the root password

```
passwd root
```
7. Force SELinux to relabel context on next boot

```
touch /.autorelabel
```
8. Remount the filesystem as read-only (not strictly necessary)

```
mount -o remount,ro /sysroot
```
9. Exit the `chroot` environment

```
exit
```
10. Resume system boot

```
exit
```

How to fix a system that is not booting to the latest installed kernel (RHEL 8 and later)

1. Reboot the system and, on the GRUB 2 boot screen, boot manually on the latest kernel X.Y.Z-AAA
2. Uninstall all kernels except the latest one


```
dnf remove kernel-core-X.Y.Z-BBB kernel-core-X.Y.Z-CCC
```
3. Remove those GRUB orphaned entries referring to the uninstalled old kernels, leaving only the ones referring to the latest kernel


```
rm /boot/loader/entries/nnnnnnn-X.Y.Z-rescue.conf
rm /boot/loader/entries/nnnnnnn-X.Y.Z-BBB.conf
rm /boot/loader/entries/nnnnnnn-X.Y.Z-rescue.conf
rm /boot/loader/entries/nnnnnnn-X.Y.Z-CCC.conf
```
4. Rebuild the GRUB configuration


```
grub2-mkconfig -o /boot/grub2/grub.cfg (for BIOS)
grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg (for UEFI)
```

How to fix a system that is not booting due to a missing UEFI partition (RHEL 9)

1. Reboot the system and, on the GRUB 2 boot screen, boot manually on another older kernel
2. Search for the modprobe configuration file that installs the VFAT module. This is usually /etc/modprobe.d/vfat.conf


```
grep vfat /etc/modprobe.d/*.conf
```
3. In the file found at the previous point, comment out the install line:


```
#install vfat /bin/true
```
4. Insert the VFAT module into the kernel and verify that it loaded correctly


```
modprobe vfat && lsmod | grep vfat
```
5. Add a line on /etc/fstab for the UEFI partition (assuming the root partition is /dev/sda1):


```
/dev/sda1 /boot/efi vfat defaults 0 1
```
6. Mount the UEFI partition


```
mount /dev/sda1 /boot/efi
```
7. Reinstall the GRUB and shim packages


```
dnf reinstall -y grub2-efi-x64 grub2-common shim-x64
```
8. Disable the auto hide menu in GRUB


```
grub2-editenv - unset menu_auto_hide
```
9. Reboot the system


```
init 6
```

How to fix a kernel panic at boot due to a bad shim signature (RHEL 9)

1. Reboot the system and, on the GRUB 2 boot screen, boot manually on another older kernel
2. Reinstall the GRUB and shim packages


```
dnf reinstall -y grub2-efi-x64 grub2-common shim-x64
```
3. Disable the auto hide menu in GRUB


```
grub2-editenv - unset menu_auto_hide
```
4. Reinstall the latest kernel


```
dnf reinstall -y kernel-X.Y.Z-AAA kernel-core-X.Y.Z-AAA
```
5. Reboot the system


```
init 6
```

How to fix a kernel panic at boot due to a missing initramfs (RHEL 9)

1. Reboot the system and, on the GRUB 2 boot screen, boot manually on another older kernel
2. Reinstall the latest kernel


```
dnf reinstall -y kernel-X.Y.Z-AAA kernel-core-X.Y.Z-AAA
```
3. Reboot the system


```
init 6
```

If the executable permission has been removed from the `chmod` command binary by mistake, any of the following procedures allows to restore it.

Copy attributes and permissions from another command binary (preserving ownership and timestamps):

```
cp --attributes-only -p /usr/bin/true /usr/bin/chmod
```

Copy the contents of `chmod` to another command binary via `cat`:

```
cp /usr/bin/true /usr/bin/true.bak
cat /usr/bin/chmod > /usr/bin/true
mv /usr/bin/true /usr/bin/chmod
mv /usr/bin/true.bak /usr/bin/true
```

Add temporarily an ACL via `setfacl`, set the executable permission, then remove the ACL:

```
setfacl -m u::rx /usr/bin/chmod
chmod +x /usr/bin/chmod
setfacl -b /usr/bin/chmod
```

Copy the binary and set permissions via `rsync`:

```
rsync /usr/bin/chmod /usr/bin/chmod2 --chmod=ugo+x
mv /usr/bin/chmod2 /usr/bin/chmod
```

Run `chmod` via the `ld` linker:

```
/usr/lib64/ld-linux-x86-64.so.2 /usr/bin/chmod +x /usr/bin/chmod (on 64-bit systems)
/usr/lib/ld-linux.so /usr/bin/chmod +x /usr/bin/chmod (on 32-bit systems)
```

Run the busybox version of `chmod`:

```
busybox chmod +x /usr/bin/chmod
```

Use the command interpreter of a programming language:

```
perl -e 'chmod 0755, "/usr/bin/chmod"' (via Perl)
python -c "import os; os.chmod('/usr/bin/chmod', 0755)" (via Python)
```

Domain Name System (DNS) is a decentralized hierarchical naming system, mostly used to resolve domain names to IP addresses. It uses TCP and UDP port 53.

DNS implementations	
BIND	Berkeley Internet Name Domain system, is the standard DNS server for UNIX
Unbound	Standard DNS server in RHEL 7
dnsmasq	Lightweight DNS, DHCP and TFTP server for a small network
djbdns	Security-hardened DNS server that also includes DNS debugging tools
PowerDNS	Alternative open source DNS server

<code>named</code>	BIND Name Daemon
<code>named -u named -g named</code>	Run BIND as user/group "named" (must be created if needed) instead of root
<code>named -t /var/cache/bind</code>	Run BIND in a chroot jail <code>/var/cache/bind</code> (it is actually the <code>chroot</code> command that starts the <code>named</code> server)
<code>ndc (UNIX)</code>	Name Daemon Control program for BIND 8
<code>rndc</code>	Remote Name Daemon Controller for BIND 9. Uses a shared key to communicate securely with <code>named</code>
<code>rndc reconfig</code>	Reload BIND configuration and new zones
<code>rndc reload example.org</code>	Reload the zone <code>example.org</code>
<code>rndc freeze example.org</code>	Suspend updates for the zone <code>example.org</code>
<code>rndc thaw example.org</code>	Resume updates for the zone <code>example.org</code>
<code>rndc tsig-list</code>	List all currently active TSIG keys
<code>dnswalk example.org.</code>	DNS debugger

DNSSEC was designed to secure the DNS tree and hence prevent cache poisoning. The TSIG (Transaction SIGNature) standard, which authenticates communications between two trusted systems, is used to sign zone transfers and DDNS (Dynamic DNS) updates.

```
dnssec-keygen -a dsa -b 1024 \  
-n HOST dns1.example.org
```

Generate a TSIG key with DNSSEC algorithm *nnn* and key fingerprint *ffff*. This will create two key files
`Kdns1.example.org.+nnn+ffff.key`
`Kdns1.example.org.+nnn+ffff.private`
which contain a key number that must be inserted both in `/etc/named.conf` and `/etc/rndc.conf`

```
rndc-confgen -a
```

Generate a `/etc/rndc.key` key file:

```
key "rndc-key" {  
    algorithm hmac-md5;  
    secret "vyZqL3tPHsqnA57e4LT0Ek==";  
};  
options {  
    default-key "rndc-key";  
    default-server 127.0.0.1;  
    default-port 953;  
};
```

This file is automatically read both by `named` and `rndc`

```
dnssec-signzone example.org
```

Sign the zone *example.org*

/etc/named.conf DNS server configuration file

```

controls {
    inet 127.0.0.1 allow {localhost;} keys {rndckey};
};
key "rndc-key" {
    algorithm dsa;
    secret "HYZur46fftdUQ43BJKI093t4t78lkp";
};

acl "mynetwork" {10.7.0.0/24;};

options {
    directory "/var/named";
    version "0.0";
    listen-on port 53 {10.7.0.1; 127.0.0.1;};
    blackhole {172.17.17.0/24;};
    allow-query {mynetwork;};
    allow-query-on {any;};
    allow-query-cache {any;};
    allow-recursion {mynetwork;};

    allow-recursion-on {mynetwork;};
    allow-transfer {10.7.0.254;};

    allow-update {any;};
    recursive-clients 1000;
    dnssec-enable yes;
    dialup no;

    forward first;
    forwarders {10.7.0.252; 10.7.0.253;};

};

// Define the root name servers
zone "." {
    type hint;
    file "root.cache";
}

// Configure system to act as a master server for the example.org domain
zone "example.org" IN {
    type master;
    file "master/example.org.zone";
};

zone "240.123.224.in-addr.arpa" IN {
    type master;
    file "slave/example.org.revzone";
};

// Configure system to act as a slave server for the example2.org domain
zone "example2.org" IN {
    type slave;
    file "slave/example2.org.zone";
    masters {10.7.0.254;};
};

zone "0.7.10.in-addr.arpa" IN {
    type slave;
    file "slave/10.7.0.revzone";
    masters {10.7.0.254;};
};

```


DNS Resource Records	
\$TTL	How long to cache a positive response
\$ORIGIN	Suffix appended to all names not ending with a dot. Useful when defining multiple subdomains inside the same zone
SOA	Start Of Authority for the example.org zone:
serial	Serial number. Must be increased after each edit of the zone file
refresh	How frequently a slave server refreshes its copy of zone data from the master
retry	How frequently a slave server retries connecting to the master
expire	How long a slave server relies on its copy of zone data. After this time period expires, the slave server is not authoritative anymore for the zone unless it can contact a master
negative TTL	How long to cache a non-existent answer
A	Address: maps names to IPv4 addresses. Used for DNS lookups.
AAAA	IPv6 address: maps names to IPv6 addresses. Used for DNS lookups.
PTR	Pointer: maps IP addresses to names. Used for reverse DNS lookups. Each A record must have a matching PTR record.
CNAME	Canonical Name: specifies an alias for a host with an A record (even in a different zone). Discouraged as it causes multiple lookups; it is better to use multiple A records instead.
NS	Name Service: specifies the authoritative name servers for the zone.
MX	Mailserver: specifies address and priority of the servers able to handle mail for the zone.
SRV	Service: specifies address and port number of the host providing a specific service. It is indicated as <code>_service._protocol.domain.</code> , where <i>protocol</i> is TCP or UDP.
Glue Records are not really part of the zone; they delegate authority for other zones, usually subdomains.	

Most common HTTP response codes		
1XX Informational	100 Continue	The server received the request headers, so the client should continue by sending the remainder of the request
	101 Switching Protocols	The server agreed to switch protocol upon client's demand
	102 Processing	The server received the request and is processing it, but response is not yet available. Used for WebDAV requests with many subrequests requiring a long time to complete; this prevents client timeout
	103 Early Hints	The client can start preloading resources while the server prepares a response. Used primarily with the Link header
2XX Success	200 OK	The request was successful
	201 Created	The request was successful, and resulted in a resource being created
	202 Accepted	The request has been received but not yet processed. Used when another server handles the request, or in batch processing
	203 Non-Authoritative Information	The returned metadata is not exactly the same as is available from the origin server. Primarily used for mirrors
	204 No Content	The request was successful, and the server does not need to return any content
	205 Reset Content	The client must reset the document which sent the request
3XX Redirection	206 Partial Content	The request was successful, and the server is returning only partial content because the client sent a Range header field
	301 Moved Permanently	The requested resource was permanently moved to a new URI
	302 Found	The requested resource was temporarily moved to a new URI
	303 See Other	The requested resource can be found on another URI, and should be retrieved from there via a GET
	304 Not Modified	The client sent a conditional GET request, and the resource has not been modified since last time it was requested
	307 Temporary Redirect	The requested resource was temporarily moved to a new URI, but future requests should use the original URI
4XX Client Error	308 Permanent Redirect	The resource is now permanently located at another URI, specified by the Location response header
	400 Bad Request	The server was unable to understand the request due to bad syntax
	401 Unauthorized	The request requires user authentication
	403 Forbidden	The client did not have the necessary permissions to access the requested resource
	404 Not Found	The requested resource was not found on the server
	408 Request Timeout	The server timed out while waiting for the request
	409 Conflict	The request could not be processed because of a conflict in the resource state
	410 Gone	The requested resource is no longer available on the server and will not be available again
5XX Server Error	451 Unavailable for Legal Reasons	The requested resource is not available due to government censorship
	500 Internal Server Error	The server encountered a generic error while trying to fulfill the request
	501 Not Implemented	The server was unable to recognize the request method
	502 Bad Gateway	The server is acting as a gateway or proxy, and received an invalid response from the upstream server
	503 Service Unavailable	The server is temporarily unavailable due to overload or maintenance
	504 Gateway Timeout	The server is acting as a gateway or proxy, and a request to the upstream server timed out
	505 HTTP Version Not Supported	The server does not support the HTTP protocol version used in the request

Apache is an open source and widespread HTTP server, originally based on the NCSA HTTPd server.

<code>apachectl</code>	(RHEL)	Manage the Apache webserver
<code>httpd</code>	(RHEL)	
<code>apache2ctl</code>	(Debian)	
<code>apachectl start</code>		Start the Apache webserver daemon
<code>apachectl status</code>		Display a brief status report
<code>apachectl fullstatus</code>		Display a detailed status report
<code>apachectl graceful</code>		Gracefully restart Apache; currently open connections are not aborted
<code>apachectl graceful-stop</code>		Gracefully stop Apache; currently open connections are not aborted
<code>apachectl configtest</code>		Test the configuration file, reporting any syntax error
<code>apachectl -t</code>		
<code>apachectl -M</code>		List all loaded and shared modules
<code>logresolve <access_log >dnslogfile</code>		Resolve IP addresses in an Apache logfile
<code>htpasswd</code>		Create or update a file used for basic HTTP authentication
<code>/etc/httpd/conf/httpd.conf</code>	(RHEL)	Apache configuration files
<code>/etc/httpd/conf.d/*.conf</code>		
<code>/etc/apache2/httpd.conf</code>	(Debian and SUSE)	
<code>/var/www/html</code>		Default document root directory
<code>\$HOME/public_html</code>		Default document root directory for users' websites
Web content must be readable by the user/group the Apache process runs as. For security reasons, it should be owned and writable by the superuser or the webmaster user/group (usually <code>www-data</code>), not the Apache user/group.		
The Apache webserver contains a number of MPMs (Multi-Processing Modules) which can operate following two methods:		
prefork MPM	A number of child processes is spawned in advance, with each child serving one connection. Highly reliable due to Linux memory protection that isolates each child process.	
worker MPM	Multiple child processes spawn multiple threads, with each thread serving one connection. More scalable but prone to deadlocks if third-party non-threadsafe modules are loaded.	

HTTPS (HyperText Transfer Protocol Secure), i.e. HTTP over SSL/TLS, allows securing communications between the webserver and the client by encrypting end-to-end all communications between the two. A webserver using HTTPS hands over its public key to the client when the client connects to the server via port 443. The server's public key is signed by a CA (Certification Authority), whose validity is ensured by the root certificates stored into the client's browser.

The commands of the OpenSSL cryptographic library (`openssl`, `CA.pl`, and `genkey`) can be used to accomplish all public key cryptography operations e.g. generate key pairs, Certificate Signing Requests, and self-signed certificates.

Virtual hosting with HTTPS requires assigning a unique IP address for each virtual host; this because the SSL handshake (during which the server sends its certificate to the client's browser) takes place before the client sends the `Host:` header (which tells to which virtual host the client wants to talk). A workaround for this is SNI (Server Name Indication) which makes the browser send the hostname in the first message of the SSL handshake. Another workaround is to have all multiple name-based virtual hosts use the same SSL certificate with a wildcard domain e.g. `*.example.org`.


Apache configuration file	
Server configuration directives	
ServerName www.mysite.org:80	Name and port (if omitted, uses default HTTP port 80) of server.
ServerRoot /etc/httpd	Root directory for configuration and log files.
ServerAdmin webmaster@mysite.org	Contact address that the server includes in any HTTP error messages to the client. Can be an email address or a URL.
StartServers 5	Number of servers to start initially.
MinSpareServers 5 MaxSpareServers 10	Minimum and maximum number of idle child server processes.
MaxClients 256 (before v2.3.13) MaxRequestWorkers 256 (v2.3.13 and later)	Max number of simultaneous requests that will be served; clients above this limit will get an HTTP error 503 - Service Unavailable. Prefork MPM: max number of child processes launched to serve requests. Worker MPM: max total number of threads available to serve requests.
ServerLimit 256	Prefork MPM: max configured value for MaxRequestWorkers. Worker MPM: in conjunction with ThreadLimit, max configured value for MaxRequestWorkers.
ThreadsPerChild 25	Worker MPM: number of threads created by each child process.
ThreadLimit 64	Worker MPM: max configured value for ThreadsPerChild.
MaxRequestsPerChild 16 (v2.2) MaxConnectionsPerChild 16 (v2.4)	Max number of connections allowed per child.
LoadModule mime_module modules/mod_mime.so	Load the module mime_module by linking in the object file or library modules/mod_mime.so.
Listen 10.17.1.1:80 Listen 10.17.1.5:8080	Make the server accept connections on the specified IP addresses (optional) and ports.
User nobody Group nobody	User and group the Apache process runs as. For security reasons, this should not be root.

Apache configuration file	
Main configuration directives	
DocumentRoot /var/www/html	Directory in filesystem that maps to the root of the website.
Alias /image /mydir/pub/image	Map the URL <code>http://www.mysite.org/image/</code> to the directory <code>/mydir/pub/image</code> in the filesystem. This allows Apache to serve content placed outside of the document root.
TypesConfig conf/mime.types	Media types file. The path is relative to <code>ServerRoot</code> .
AddType image/jpeg jpeg jpg jpe	Map the specified filename extensions onto the specified content type. These entries add to or override the entries from the media types file <code>conf/mime.types</code> .
Redirect permanent /foo /bar	Redirect to a URL on the same host. Status can be: permanent Return an HTTP status "301 - Moved Permanently" temp Return an HTTP status "302 - Found" (default) seeother Return an HTTP status "303 - See Other" gone Return an HTTP status "410 - Gone"
Redirect /foo http://www.example.com/foo	Redirect to a URL on a different host.
AccessFileName .htaccess	Name of the distributed configuration file, which contains directives that apply to the document directory it is in and to all its subtrees.
<pre><Directory "/var/www/html/foobar"> AllowOverride AuthConfig Limit </Directory></pre>	Specify which global directives an <code>.htaccess</code> file can override: AuthConfig Authorization directives for directory protection FileInfo Document type and metadata Indexes Directory indexing Limit Host access control Options Specific directory features All All directives None No directive
Limited scope directives	
<pre><Directory "/var/www/html/foobar"> [list of directives] </Directory></pre>	Limit the scope of the specified directives to the directory <code>/var/www/html/foobar</code> and its subdirectories.
<pre><Location /foobar> [list of directives] </Location></pre>	Limit the scope of the specified directive to the URL <code>http://www.mysite.org/foobar/</code> and its subdirectories.
Logging directives	
LogFormat "%h %l %u %t \"%r\" %>s %b"	Specify the format of a log.
LogFormat "%h %l %u %t \"%r\" %>s %b" common	Specify a nickname for a log format. In this case, specifies "common" for the CLF (Common Log Format) which is defined as such: %h IP address of the client host %l Identity of client as determined by <code>identd</code> %u User ID of client making the request %t Timestamp the server completed the request %r Request as done by the user %s Status code sent by the server to the client %b Size of the object returned, in bytes
CustomLog /var/log/httpd/access_log common	Set up a log filename, with the format or (as in this case) the nickname specified.
TransferLog /var/log/httpd/access_log	Set up a log filename, with format determined by the most recent <code>LogFormat</code> directive which did not define a nickname.
TransferLog " rotatelog access_log 86400"	Set log rotation every 24 hours.
HostnameLookups Off	Disable DNS hostname lookup to save network traffic. Hostnames can be resolved later by post-processing the log file via the <code>logresolve</code> command.

Apache configuration file	
Virtual hosts directives	
NameVirtualHost * (v2.2)	Specify which IP address will serve virtual hosting. The argument can be an IP address, an <i>address:port</i> pair, or * for all IP addresses of the server. The same argument need to be inserted in the relevant <VirtualHost> directive.
<VirtualHost *:80> ServerName www.mysite.org ServerAlias mysite.org *.mysite.org DocumentRoot /var/www/vhosts/mysite </VirtualHost>	The first listed virtual host is also the default virtual host. It inherits those main settings that does not override. This virtual host answers to http://www.mysite.org , and also redirects there all HTTP requests on the domain mysite.org.
<VirtualHost *:80> ServerAdmin webmaster@www.mysite2.org ServerName www.mysite2.org DocumentRoot /var/www/vhosts/mysite2 ErrorLog /var/www/logs/mysite2 </VirtualHost>	Name-based virtual host http://www.mysite2.org . Multiple name-based virtual hosts can share the same IP address; DNS must be configured accordingly to map each name to the correct IP address. Cannot be used with HTTPS.
<VirtualHost *:8080> ServerName www.mysite3.org DocumentRoot /var/www/vhosts/mysite3 </VirtualHost>	Port-based virtual host answering to connections on port 8080. A <code>Listen 8080</code> directive must also be present.
<VirtualHost 10.17.1.5:80> ServerName www.mysite4.org DocumentRoot /var/www/vhosts/mysite4 </VirtualHost>	IP-based virtual host answering to http://10.17.1.5 .


Apache configuration file	
Authorization directives	
AuthName "Protected zone"	Name of the realm. The client will be shown the realm name and prompted to enter a user and password.
AuthType Basic	Type of user authentication: Basic, Digest, Form, or None.
AuthUserFile "/var/www/.htpasswd"	User database file. Each line has the format <i>user:encryptedpassword</i> . To add a user to the database file, use the command: <code>htpasswd /var/www/.htpasswd user</code> (will prompt for a password).
AuthGroupFile "/var/www/.htgroup"	Group database file. Each line specifies a group followed by the usernames of all its members: <i>group: user1 user2 user3</i>
Require valid-user	Control who can access the protected resource. valid-user Any user in the user database file user user Only the specified user group group Only the members of the specified group
Satisfy Any	Set the access policy concerning user and host control. All Both Require and Allow criteria must be satisfied Any Any of Require or Allow criteria must be satisfied
Allow from 10.13.13.0/24 Deny from 10.13.14.0/24 (deprecated)	Control which host can or cannot access the protected resource.
Order Allow,Deny (deprecated)	Control the evaluation order of Allow and Deny directives. Allow,Deny First, all Allow directives are evaluated; at least one must match, or the request is rejected. Next, all Deny directives are evaluated; if any matches, the request is rejected. Last, any requests which do not match an Allow or a Deny directive are denied. Deny,Allow First, all Deny directives are evaluated; if any match, the request is denied unless it also matches an Allow directive. Any requests which do not match any Allow or Deny directives are permitted.

Apache configuration file	
SSL/TLS directives (<code>mod_ssl</code> module)	
<code>SSLCertificateFile \</code> <code>/etc/httpd/conf/ssl.crt/server.crt</code>	SSL server certificate.
<code>SSLCertificateKeyFile \</code> <code>/etc/httpd/conf/ssl.key/server.key</code>	SSL server private key (for security reasons, this file must be mode 600 and owned by root).
<code>SSLCACertificatePath \</code> <code>/usr/local/apache2/conf/ssl.crt/</code>	Directory containing the certificates of CAs. Files in this directory are PEM-encoded and accessed via symlinks to hash filenames.
<code>SSLCACertificateFile \</code> <code>/usr/local/apache2/conf/ssl.crt/ca-bundle.crt</code>	Certificates of CAs. Certificates are PEM-encoded and concatenated in a single bundle file in order of preference.
<code>SSLCertificateChainFile \</code> <code>/usr/local/apache2/conf/ssl.crt/ca.crt</code>	Certificate chain of the CAs. Certificates are PEM-encoded and concatenated from the issuing CA certificate of the server certificate to the root CA certificate. Optional.
<code>SSLEngine on</code>	Enable the SSL/TLS Protocol Engine.
<code>SSLProtocol +SSLv3 +TLSv1.2</code>	SSL protocol flavors that the client can use to connect to the server. Possible values are: SSLv2 (deprecated) SSLv3 TLSv1 TLSv1.1 TLSv1.2 All (all the above protocols)
<code>SSLCipherSuite \</code> <code>ALL:!aDH:RC4+RSA:+HIGH:+MEDIUM:+LOW:+SSLv2:+EXP</code>	Cipher suite available for the SSL handshake (key exchange algorithms, authentication algorithms, cipher/encryption algorithms, MAC digest algorithms).
<code>ServerTokens Full</code>	Server response header field to send back to client. Possible values are: Prod sends Server: Apache Major sends Server: Apache/2 Minor sends Server: Apache/2.4 Minimal sends Server: Apache/2.4.2 OS sends Server: Apache/2.4.2 (Unix) Full sends Server: Apache/2.4.2 (Unix) \ PHP/4.2.2 MyMod/1.2 (default)
<code>ServerSignature Off</code>	Trailing footer line on server-generated documents. Possible values are: Off no footer line (default) On server version number and <code>ServerName</code> EMail as above, plus a mailto link to <code>ServerAdmin</code>
<code>SSLVerifyClient none</code>	Certificate verification level for client authentication. Possible values are: none no client certificate is required require the client needs to present a valid certificate optional the client may present a valid certificate (this option is unused as it doesn't work on all browsers) optional_no_ca the client may present a valid certificate but it doesn't need to be successfully verifiable (this option is practically used only for SSL testing)
<code>TraceEnable on</code>	Enable TRACE requests.



A **forward proxy** provides proxy services, typically web content caching and/or filtering, for clients located in a LAN. All outgoing requests from the clients, and the responses from the Internet, pass through the proxy. The clients must be manually configured (e.g. in the browser's connection settings) to use the proxy.

Apache configuration file	
Forward proxy	
ProxyRequests On	Enable forward proxy requests.
ProxyVia On	Add a <code>Via: HTTP</code> header line to every request and reply.
<pre><Proxy "*"> Require ip 10.1.1 </Proxy></pre>	Serve only proxy requests coming from 10.1.1.0/24.



A **reverse proxy** aka **gateway** allows to expose a single entry point for one or more webservers in a LAN. This improves security and simplifies management, as features (e.g. load balancing, firewalling, automatic redirection from HTTP to HTTPS, redirection on default ports) can be configured centrally. It is necessary to create a DNS A record that maps `site.example.com` to the public IP address of the proxy.

Apache configuration file	
Reverse proxy	
<VirtualHost *:80>	Virtual host for HTTP.
ServerName site.example.com	Define website name.
RewriteEngine On RewriteCond %{HTTPS} off RewriteRule (.*) https://%{HTTP_HOST}%{REQUEST_URI}	Redirect all HTTP requests to HTTPS.
Alternatively: Redirect "/" "https://10.2.2.73:443/"	
</VirtualHost>	
<VirtualHost *:443>	Virtual host for HTTPS.
ServerName site.example.com	Define website name.
ServerSignature On	Set a footer line under server-generated pages.
<Proxy *> Require all granted </Proxy>	Serve all proxy requests.
SSLEngine on SSLProtocol ALL -SSLv2 -SSLv3 SSLHonorCipherOrder on SSLCipherSuite DEFAULT SSLCertificateFile /etc/httpd/ssl/site.crt SSLCertificateKeyFile /etc/httpd/ssl/site.key SSLCACertificateFile /etc/httpd/ssl/site.ca.crt	Enable and configure SSL.
ProxyPass "/" "http://10.2.2.73:8080/" ProxyPassReverse "/" "http://10.2.2.73:8080/"	Enable reverse proxying for server 10.2.2.73.
</VirtualHost>	

<code>nginx</code>	HTTP and reverse proxy server. Stable, high performing, and requiring low levels of resources
<code>thttpd</code>	Tiny / turbo / throttling HTTP server. Small and fast webserver
<code>lighttpd</code>	Fast and flexible HTTP server, optimized for high performance environments

Apache Tomcat is an open source Java Servlet Container implementing several Java EE specifications, originally part of the Jakarta Project. It is composed of:

- Catalina, the core component and servlet container implementation;
- Coyote, an HTTP connector component, providing a pure Java webserver environment to run Java code;
- Jasper, a JSP (Java Server Pages) engine, which parses JSP files and compiles them into Java servlets.

Tomcat has been removed from RHEL 8 and later; instead, it is suggested to use the JBoss Enterprise Application Platform, which includes Apache and Tomcat.

<code>\$JAVA_HOME</code>	Root of the Java installation e.g. <code>/usr/lib/jvm/java-1.8.0-openjdk.x86_64/</code>
<code>\$CATALINA_HOME</code>	Root of the Tomcat installation e.g. <code>/usr/share/tomcat7/</code>
<code>\$CATALINA_BASE</code>	Tomcat may also be configured for multiple instances by defining the variable <code>\$CATALINA_BASE</code> for each instance. If a single instance of Tomcat is running, <code>\$CATALINA_BASE</code> is the same as <code>\$CATALINA_HOME</code>

Tomcat global files	
<code>\$CATALINA_BASE/conf/server.xml</code>	Tomcat main configuration file
<code>\$CATALINA_BASE/conf/web.xml</code>	Options and values applied to all web applications running on a specific Tomcat instance. These can be overridden by the application-specific servlet configuration defined in <code>\$CATALINA_BASE/webapps/appname/WEB-INF/web.xml</code>
<code>\$CATALINA_BASE/conf/context.xml</code>	Context applied to all web applications running on a specific Tomcat instance
<code>\$CATALINA_BASE/conf/tomcat-users.xml</code>	Users, passwords, and roles applied to a specific Tomcat instance
<code>\$CATALINA_BASE/conf/catalina.policy</code>	Tomcat's core security policy for the Catalina class
<code>\$CATALINA_BASE/conf/catalina.properties</code>	Java properties file for the Catalina class
<code>\$CATALINA_BASE/conf/logging.properties</code>	Java properties file for Catalina's built-in logging functions
<code>\$CATALINA_BASE/lib/</code>	JAR files accessible by both web applications and internal Tomcat code
<code>\$JAVA_HOME/jre/lib/security/keystore.jks</code>	Java keystore
Tomcat application-specific files	
<code>\$CATALINA_BASE/webapps/appname/WEB-INF/</code>	HTML, JSP, and other files to serve to the client browser
<code>\$CATALINA_BASE/webapps/appname/WEB-INF/web.xml</code>	Description of servlets and other components of the application, and initialization parameters
<code>\$CATALINA_BASE/webapps/appname/WEB-INF/classes/</code>	Java class files that aren't in JAR format. The directory hierarchy from here reflects the class hierarchy
<code>\$CATALINA_BASE/webapps/appname/WEB-INF/lib/</code>	Other JAR files (e.g. third-party libraries, JDBC drivers) required by the application
Tomcat log files	
<code>\$CATALINA_BASE/logs/catalina.out</code>	Tomcat log
<code>\$CATALINA_BASE/logs/localhost.log</code>	Host log
<code>\$CATALINA_BASE/logs/localhost_access.log</code>	Host HTTP access log
<code>\$CATALINA_BASE/logs/manager.log</code>	Application log
<code>\$CATALINA_BASE/logs/host-manager.log</code>	Application log

<code>java -X</code>	Display all available <code>-x</code> options (nonstandard HotSpot JVM options)
<code>java -XshowSettings:properties -version</code>	Print Java runtime settings

Samba is a free-software, cross-platform implementation of SMB/CIFS. **SMB (Server Message Block)** is a Microsoft proprietary protocol for file and printer sharing, while **CIFS (Common Internet File System)** is the public version of SMB.

Commonly used ports in Samba		
TCP/UDP 137	netbios-ns	NetBIOS Name Service requests and responses
TCP/UDP 138	netbios-dgm	NetBIOS Datagram Service e.g. server announcements
TCP/UDP 139	netbios-ssn	NetBIOS Session Service e.g. file and printer sharing
TCP 445	microsoft-ds	Active Directory; registration and translation of NetBIOS names, network browsing
TCP 389		LDAP
TCP 901		SWAT service

The full list of used ports can be found via the command `grep -i netbios /etc/services`.

<code>smbd</code>	Server Message Block daemon. Provides SMB file and printer sharing, browser services, user authentication, and resource lock. An extra copy of this daemon runs for each client connected to the server
<code>nmbd</code>	NetBIOS Name Service daemon. Handles NetBIOS name lookups, WINS requests, list browsing and elections. An extra copy of this daemon runs if Samba functions as a WINS server; another extra copy of this daemon runs if DNS is used to translate NetBIOS names. WINS (Windows Internet Name Service) is a name service used to translate NetBIOS names to IP addresses
<code>/etc/smb/</code>	Samba directory
<code>/etc/samba/</code> (RHEL 7)	
<code>/etc/samba/lmhosts</code>	Samba NetBIOS hosts file
<code>/etc/samba/netlogon</code>	User logon directory
<code>smbd -V</code> <code>smbclient -V</code>	Show the version of the Samba server
<code>testparm</code>	Check the Samba configuration file and report any error
<code>smbpasswd user</code>	Change the Samba password of <i>user</i>
<code>smbpasswd -a user</code>	Create a new Samba <i>user</i> and set their password
<code>nmblookup smbserver</code>	Look up the NetBIOS name of a server and map it to an IP address
<code>nmblookup -U winsserver -R WORKGROUP#1B</code>	Query recursively a WINS server for the Domain Master Browser for the specified workgroup
<code>nmblookup -U winsserver -R WORKGROUP#1D</code>	Query recursively a WINS server for the Domain Controller for the specified workgroup
<code>net</code>	Tool for administration of Samba and remote CIFS servers
<code>net rpc shutdown -r -S smbserver -U root%password</code>	Reboot a CIFS server
<code>net rpc service list -S smbserver</code>	List available services on a CIFS server
<code>net status sessions</code>	Show active Samba sessions
<code>net status shares</code>	Show Samba shares
<code>net rpc info</code>	Show information about the domain
<code>net groupmap list</code>	Show group mappings between Samba and MS Windows

```
mount.cifs  
smbmount
```

Mount a Samba share on a Linux filesystem, using the CIFS filesystem interface

```
mount //smbserver/share1 /mnt/share1 \  
-t cifs -o username=user
```

Mount a Samba share as *user*

```
smbstatus
```

Display current information about shares, clients connections, and locked files

```
smbclient //smbserver/share1  
smbclient -L //smbserver -W WORKGROUP -U user
```

Access a Samba share on a server (with an FTP-like interface)

List the Samba resources available on a server, belonging to the specified workgroup and accessible to the specified user

```
cat msg.txt | smbclient -M client -U user
```

Show a message popup on the client machine, using the WinPopup protocol

Samba mount options	
<code>username=user</code>	Mount the share as <i>user</i>
<code>password=password</code>	Specify the mount user's <i>password</i>
<code>credentials=credfile</code>	Mount the share as the user defined in the credentials file <i>credfile</i> which must have this format: <code>username=user</code> <code>password=password</code>
<code>multiuser</code>	Mount the share in multiuser mode
<code>sec=ntlmssp</code>	Set the security level to NTLMSSP. This is required in RHEL 7 to enable multiuser mode

/etc/samba/smb.conf Samba configuration	
<pre>[global] workgroup = MYWORKGROUP server string = Linux Samba Server %L hosts allow = 10.9.9.0/255.255.255.0 security = user encrypt passwords = yes smb passwd file = /etc/samba/smbpasswd unix password sync = yes username map = /etc/samba/smbusers netbios name = Mysambabox netbios aliases = Mysambabox1 wins support = yes logon server = yes log file = /var/log/samba/log.%m max log size = 1000 syslog only = no syslog = 0 panic action = \ /usr/share/samba/panic-action %d</pre>	Global server settings: defines parameters applicable for the whole Samba server and sets the defaults that will be used for the parameters not mentioned in other sections Make Samba join the specified workgroup Describe server to the clients Allow only the specified machines to connect to the server Set up user-level authentication Use encrypted passwords Refer to the specified password file for user authentication. A new user's password will need to be set both in Linux and Samba by running these commands from a shell prompt: <pre>passwd newuser smbpasswd newuser</pre> When the password of a client user (e.g. under MS Windows) is changed, change the Linux and Samba passwords accordingly Map each Samba server user name to client user name(s). The file <code>/etc/samba/smbusers</code> has the following format: <pre>root = Administrator Admin jdoe = "John Doe" kgreen = "Kim Green"</pre> Set NetBIOS name and alias Make Samba play the role of a WINS server. Note: There should be only one WINS server on a network Enable logon support. Logon script parameters will be defined in a <code>[netlogon]</code> section Use a separate logfile for each machine that connects Maximum size of each logfile, in Kb Do not use only syslog to log Log everything to the logfiles <code>/var/log/smb/log.smbd</code> and <code>/var/log/smb/log.nmbd</code>, and log a minimum amount of information to syslog. This parameter can be set to a higher value to have syslog log more information Mail a backtrace to the sysadmin in case Samba crashes
<pre>[netlogon] comment = Netlogon for Windows clients path = /home/netlogon logon script = %U.bat browseable = no writeable = no guest ok = no</pre>	Section defining a logon script Specifies a per-user script e.g. <code>/home/netlogon/jdoe.bat</code> will be called when user <code>jdoe</code> logs in. It is also possible to specify a per-clientname script <code>%m.bat</code>, which will be called when a specific machine logs in. Guest access to the service (i.e. access without entering a password) is disabled
<pre>[Canon LaserJet 3] printer name = lp comment = Canon LaserJet 3 main printer path = /var/spool/lpd/samba printable = yes writeable = no</pre>	Section defining a printer accessible via the network

/etc/samba/smb.conf Samba configuration	
<pre>[public] comment = Public Storage on %L path = /home/samba browsable = yes writeable = yes</pre>	Section defining a public share accessible on read/write by anyone Describe the public share to users Path of the public share on the server Show the public share when browsing Allow all users to write in this directory
<pre>[homes] comment = %U's home directory on %L from %m browseable = no writeable = yes</pre>	Section enabling users that have an account and a home directory on the Samba server to access it and modify its contents from a Samba client. The <code>path</code> variable is not set, by default is <code>path=/home/%S</code> Describe the share to the user Do not show the homes share when browsing Allow the user to write in their home directory
<pre>[foobar] path = /foobar comment = Share Foobar on %L from %m browsable = yes writeable = yes valid users = jdoe, kgreen, +geeks invalid users = csmith read list = bcameron write list = fcastle</pre>	Section defining a specific share Path of the share on the server Describe the share to users Show the share when browsing Allow the users to write in this share Allow access only to users "jdoe" and "kgreen", and to local group "geeks" Deny access to user "csmith" Allow read-only access to user "bcameron" Allow read-write access to user "fcastle"

/etc/samba/smb.conf Samba configuration	
User-level authentication	
<pre>[global] security = user guest account = nobody map to guest = Never</pre>	Set up user-level authentication Map the guest account to the system user nobody (default) Specify how incoming requests are mapped to the guest account: Bad User redirect from an invalid user to guest account on server Bad Password redirect from an invalid password to guest account on server Never reject unauthenticated users
Server-level authentication	
<pre>[global] security = server password server = srv1 srv2</pre>	Set up server-level authentication Authenticate to server <i>srv1</i>, or to server <i>srv2</i> if the first one is unavailable
Domain-level authentication	
<pre>[global] security = ADS realm = KRB_REALM</pre>	Set up domain-level authentication as an Active Directory member server Join the specified realm. Kerberos must be installed and an administrator account must be created: <pre>net ads join -U Administrator%password</pre>
Share-level authentication	
<pre>[global] security = share [foobar] path = /foobar username = user only user = yes</pre>	Set up share-level authentication Define a "foobar" share accessible to any user which can supply <i>user's</i> password. The <i>user</i> must be created on the system: <pre>useradd -c "Foobar account" -d /tmp -m -s /sbin/nologin user</pre> and added to the Samba password file: <pre>smbpasswd -a user</pre>

Samba macros		
%S	Username	These macros are applied only to configuration options used once a connection has been established:
%U	Session username i.e. the username that the client requested, not necessarily the same as the one the client obtained	
%G	Primary group of session username	%S Name of the current service, if any
%h	Samba server hostname	%P Root directory of the current service, if any
%M	Client hostname	%u Username of the current service, if any
%L	NetBIOS name of the server	%g Primary group name of username
%m	NetBIOS name of the client	%H Home directory of username
%d	Process ID of the current server process	%N Name of the NIS home directory server as obtained from the NIS <code>auto.map</code> entry. Same as %L if Samba was not compiled with the <code>--with-automount</code> option
%a	Architecture of remote machine	%p Path of service's home directory as obtained from the NIS <code>auto.map</code> entry. The NIS <code>auto.map</code> entry is split up as %N:%p
%I	IP address of client machine	
%i	Local IP address to which a client connected	
%T	Current date and time	
%D	Domain or workgroup of the current user	
%w	Winbind separator	
%(var)	Value of the environment variable <i>var</i>	

Samba setup

This procedure allows sharing on read-write the local directory `/smbshare` on server 10.1.1.1 to client 10.2.2.2.

Server setup:

1. Create the group for write access to the share `groupadd -r geeks`
2. Create the user and assign it to the group `useradd -G geeks jdoe`
3. Add the user to Samba.
You will be prompted to enter a password `smbpasswd -a jdoe`
4. Assign correct ownership to the share `chgrp geeks /smbshare`
5. Set the SGID bit to the share `chmod 2775 /smbshare`
6. Set the correct SELinux label to the share `semanage fcontext -a -t samba_share_t '/smbshare'`
`restorecon -FR /smbshare`
7. Enable the SELinux boolean for write access to the share `setsebool -P samba_export_all_rw=on`
8. Add a section for the share on `/etc/samba/smb.conf`:


```
[smbshare]
  path = /smbshare
  hosts allow = 10.2.2.2
  write list = @geeks
```
9. Ensure that the `smb` and `nmb` services are running

Client setup:

1. Add an entry to `/etc/fstab` to mount the Samba share device automatically:


```
//10.1.1.1/smbshare /mountpoint cifs username=jdoe,password=s3cr3t 0 0
```

Client multiuser setup:

1. Add an entry to `/etc/fstab` to mount the Samba share device automatically in multiuser mode:


```
//10.1.1.1/smbshare /mountpoint cifs username=jdoe,password=s3cr3t,multiuser,sec=ntlmssp 0 0
```
2. Login as another user (there must be a matching Samba user on the Samba server 10.1.1.1) `su - ksmith`
3. Store the Samba username and password in the kernel keyring for the current session `cifscreds add 10.1.1.1`

A **Network File System (NFS)** server makes filesystems available to remote clients for mounting.

NFS requires the portmapper to map incoming TCP/IP connections to the appropriate NFS RPC calls. Some Linux distributions use rpcbind instead of the portmapper.

For security reasons, the TCP Wrapper should be configured to limit access to the portmapper to NFS clients only:

file /etc/hosts.deny should contain portmap: ALL

file /etc/hosts.allow should contain portmap: IP_addresses_of_clients

NFS handles user permissions across systems by considering users with same UID and username as the same user. Group permission is evaluated similarly, by GID and groupname.

rpc.nfsd
rpc.mountd
rpc.lockd
rpc.statd

NFS daemons

/etc/exports

List of the filesystems to be exported (via the command `exportfs`)

/var/lib/nfs/xtab

List of exported filesystems, maintained by `exportfs`

/proc/fs/nfs/exports

Kernel export table (can be examined via the command `cat`)

`exportfs -ra`

Export or reexport all directories.
When exporting, fills the kernel export table `/proc/fs/nfs/exports`.
When reexporting, removes the entries in `/var/lib/nfs/xtab` that are deleted from `/etc/exports` (therefore synchronizing the two files), and removes the entries from `/proc/fs/nfs/exports` that are no longer valid

`exportfs -ua`

Unexport all directories.
Removes from `/proc/fs/nfs/exports` the entries that are listed in `/var/lib/nfs/xtab`, and clears the latter file

`mount -t nfs nfsserver:/share /usr`

Command to be run on a client to mount locally a remote NFS share. NFS shares accessed frequently should be added to `/etc/fstab` e.g.
`nfsserver:/share /usr nfs intr 0 0`

`showmount`

Show the remote client hosts currently having active mounts

`showmount --directories`

Show the directories currently mounted by a remote client host

`showmount --exports`

Show the filesystems currently exported i.e. the active export list

`showmount --all`

Show both remote client hosts and directories

`showmount -e nfsserver`

Show the shares a NFS server has available for mounting

`rpcinfo -p nfsserver`

Probe the portmapper on a NFS server and display the list of all registered RPC services there

`rpcinfo -t nfsserver nfs`

Test a NFS connection by sending a null pseudo request (using TCP)

`rpcinfo -u nfsserver nfs`

Test a NFS connection by sending a null pseudo request (using UDP)

`nfsstat`

Display NFS/RPC client/server statistics.

Options:

	NFS	RPC	both
server	-sn	-sr	-s
client	-cn	-cr	-c
both	-n	-r	-nr

/etc/exports	
/export/	10.3.3.3(rw)
/export2/	10.4.4.0/24
/export3/	*(ro,sync)
/home/ftp/pub	myhost(rw) *.example.org(ro)
/home/crew	@FOOWORKGROUP(rw) (ro)

filesystem	Filesystem on the NFS server to be exported to clients	
client identity	Client systems permitted to access the exported directory. Can be specified by hostname, IP address, wildcard, subnet, or @NIS workgroup. Multiple client systems can be listed, and each one can have different options	
client options	ro	Read-only access (default)
	rw	Read and write access. The client might choose to mount read-only anyway
	sync	Reply to requests only after the changes made by these requests have been committed to stable storage
	async	Reply to requests without waiting that changes are committed to stable storage. Improves performances but might cause loss or corruption of data if server crashes
	root_squash	Requests by user <code>root</code> on client will be done as user <code>nobody</code> on server (default)
	no_root_squash	Requests by user <code>root</code> on client will be done as same user <code>root</code> on server
	all_squash	Requests by a non-root user on client will be done as user <code>nobody</code> on server
	no_all_squash	Requests by a non-root user on client will be attempted as same user on server (default)

NFS mount options	
rsiz=nnn	Size for read transfers (from server to client)
wsiz=nnn	Size for write transfers (from client to server)
nfsvers=n	Use NFS version <i>n</i> for transport
retry=n	Keep retrying a mount attempt for <i>n</i> minutes before giving up
timeo=n	A mount attempt times out after <i>n</i> tenths of a second
intr	User can interrupt a mount attempt
nointr	User cannot interrupt a mount attempt (default)
hard	The system will try a mount indefinitely (default)
soft	The system will try a mount until an RPC timeout occurs
bg	Try a mount in the foreground; all retries occur in the background
fg	All mount attempts occur in the foreground (default)
tcp	Connect using TCP
udp	Connect using UDP
sec=krb5p	Use Kerberos to encrypt all requests between client and server
v4.2	Enable NFS v4.2, which allows the server to export the SELinux context

NFS setup

This procedure allows sharing on read-write mode the local directory `/nfsshare` on server 10.1.1.1 to client 10.2.2.2.

Server setup:

1. Ensure that the `nfs-server` service is running
2. Change ownership of the share `chown nfsnobody /nfsshare`
3. Add an entry for the share on `/etc/exports`:
`/nfsshare 10.2.2.2(rw)`
4. Reload the exports file `exportfs -r`

Client setup:

1. Add an entry to `/etc/fstab` to mount the NFS share device automatically:
`10.1.1.1:/nfsshare /mountpoint nfs defaults 0 0`

Secure NFS setup

This procedure allows sharing on read-write mode the local directory `/nfsshare` on server 10.1.1.1 to client 10.2.2.2, securely with Kerberos enabled.

Server setup:

1. Install the appropriate server keytab on `/etc/krb5.keytab`
2. Ensure that the `nfs-secure-server` service is running
3. Change ownership of the share `chown nfsnobody /nfsshare`
4. Add an entry for the share on `/etc/exports`:
`/nfsshare 10.2.2.2(sec=krb5p,rw)`
5. Reload the exports file `exportfs -r`

Client setup:

1. Install the appropriate client keytab on `/etc/krb5.keytab`
2. Ensure that the `nfs-secure` service is running
3. Add an entry to `/etc/fstab` to mount the NFS share device automatically:
`10.1.1.1:/nfsshare /mountpoint nfs defaults,sec=krb5p 0 0`

iSCSI (Internet Small Computer System Interface) is a network protocol that allows emulating an SCSI local storage device over a TCP/IP network. By default it uses TCP port 3260.

An iSCSI server can use a local block device (physical or virtual disk, disk partition, or Logical Volume), a file, a physical SCSI device, or a ramdisk as the underlying storage resource (**backstore**) and make it available by assigning it a **LUN** (Logical Unit Number). An iSCSI server provides one or more **targets**, each of which presents one or more LUNs and is able to accept connections from an iSCSI client (**initiator**).

Targets and initiators are called **nodes** and are identified by a unique **IQN** (iSCSI Qualified Name) e.g.

`iqn.2017-11.org.example.subdomain:foo:bar`. The IP address and port of a node is called a **portal**.

A target accepts connections from an initiator via a **TPG** (Target Portal Group) i.e. its IP address and port. A TPG may have an ACL in place so to accept connections only from a specific initiator's IQN.

`targetcli` Target configurator (server side). Can be used as a command line tool or as an interactive shell. Configuration is saved to `/etc/target/saveconfig.json`

`iscsiadm` Administration tool for iSCSI devices (client side)

iSCSI setup

This procedure makes available the local disk `/dev/sdb` on server 10.1.1.1 to the client having IQN `iqn.2017-11.org.example:client`.

Server (target) setup:

1. Ensure that the `targetcli` service is running
2. Enter the `targetcli` shell


```
targetcli
```
3. Create a backstore


```
cd /backstores/block
create mydisk /dev/sdb
```
4. Create a IQN for the target.
This automatically creates a TPG for the IQN


```
cd /iscsi
create iqn.2017-11.org.example:target
```
5. On the TPG, create an ACL to allow connections from the initiator with a specific IQN


```
cd /iscsi/iqn.2017-11.org.example:target/tpg1/acls
create iqn.2017-11.org.example:client
```
6. On the TPG, create a LUN for the backstore


```
cd /iscsi/iqn.2017-11.org.example:target/tpg1/luns
create /backstores/block/mydisk
```
7. On the TPG, create a portal listening from the server's IP address


```
cd /iscsi/iqn.2017-11.org.example:target/tpg1/portals
delete 0.0.0.0 ip_port=3260
create 10.1.1.1
```
8. Verify the configuration


```
ls /
```

```
o- / ..... [...]
  o- backstores ..... [Storage Objects: 1]
    | o- block ..... [Storage Objects: 1]
    | | o- mydisk ..... [/dev/sdb (100.0MiB) write-thru activated]
    | |   o- alua ..... [ALUA Groups: 1]
    | |     o- default_tg_pt_gp ..... [ALUA state: Active/optimized]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 1]
    | o- iqn.2017-11.org.example:target ..... [TPGs: 1]
    |   o- tpg1 ..... [no-gen-acls, no-auth]
    |     o- acls ..... [ACLs: 1]
    |       | o- iqn.2017-11.org.example:client ..... [Mapped LUNs: 1]
    |       |   o- mapped_lun0 ..... [lun0 block/mydisk (rw)]
    |       o- luns ..... [LUNs: 1]
    |         | o- lun0 ..... [block/mydisk (/dev/sdb) (default_tg_pt_gp)]
    |         o- portals ..... [Portals: 1]
    |           o- 10.1.1.1:3260 ..... [OK]
    o- loopback ..... [Targets: 0]
```
9. Exit the `targetcli` shell.


```
exit
```

Configuration is automatically saved

Client (initiator) setup:

1. Set the correct initiator IQN in the file `/etc/iscsi/initiatorname.iscsi`:


```
InitiatorName=iqn.2017-11.org.example:client
```
2. Ensure that the `iscsi` service is running
3. Discover the iSCSI target(s) provided by the portal. This echoes the target(s) IQN found


```
iscsiadm -m discovery -t sendtargets -p 10.1.1.1
```
4. Login to the target IQN found


```
iscsiadm -m node -T iqn.2017-11.org.example:target -p 10.1.1.1 -l
```

The iSCSI device is now locally available and can be formatted and mounted. Node records remain after logout or reboot; the system will login again to the target IQN automatically
5. Add an entry to `/etc/fstab` to mount the iSCSI device automatically:


```
UUID=nnnnnnnnn-nnnn-nnnn-nnnn-nnnnnnnnnnnn /mountpoint fstype _netdev 0 0
```

DHCP (Dynamic Host Configuration Protocol) is a protocol for network management that automatically assigns to a requesting host an IP address and other network configuration parameters. It is based on **BOOTP (Bootstrap Protocol)**. A DHCP server listens for requests on UDP port 67 and answers to UDP port 68. The assignment of an IP address to a host is done through a sequence of DHCP messages initiated by the client host, which is (for DHCPv4): DHCP Discover, DHCP Offer, DHCP Request, and then DHCP ACK.

Because DHCP Discover messages are broadcast and therefore not routed outside a LAN, a DHCP relay agent is necessary for those clients situated outside the DHCP server's LAN. The DHCP relay agent listens to DHCP Discover messages and relays them in unicast to the DHCP server.

<code>/etc/dhcpd.conf</code>	Configuration file for the DHCP server
<code>/etc/sysconfig/dhcrelay</code> (SUSE)	Configuration file for the DHCP relay agent
<code>/var/lib/dhcpd/dhcpd.leases</code>	DHCP current leases

<code>/etc/dhcpd.conf</code>	DHCP server configuration
<pre>option domain-name-servers 10.2.2.2; option smtp-servers 10.3.3.3; option pop-servers 10.4.4.4; option time-servers 10.5.5.5; option nntp-servers 10.6.6.6;</pre>	Global parameters for DNS, mail, NTP, and news servers specification
<pre>shared-network geek-net { default-lease-time 86400; max-lease-time 172800; option routers 10.0.3.252; option broadcast-address 10.0.3.255; subnet 10.0.3.0 netmask 255.255.255.128 { range 10.0.3.1 10.0.3.101; } subnet 10.0.3.128 netmask 255.255.255.128 { range 10.0.3.129 10.0.3.229; } }</pre>	Definition of a network Time, in seconds, that will be assigned to a lease if a client does not ask for a specific expiration time Maximum time, in seconds, that can be assigned to a lease if a client asks for a specific expiration time Definition of different subnets in the network, with specification of different ranges of IP addresses that will be leased to clients depending on the client's subnet
<pre>group { option routers 10.0.17.252; option broadcast-address 10.0.17.255; netmask 255.255.255.0; host linuxbox1 { hardware ethernet AA:BB:CC:DD:EE:FF; fixed-address 10.0.17.42; option host-name "linuxbox1"; } host linuxbox2 { hardware ethernet 33:44:55:66:77:88; fixed-address 10.0.17.66; option host-name "linuxbox2"; } }</pre>	Definition of a group Definition of different hosts to whom static IP addresses will be assigned to, depending on their MAC address

DHCPv4 message	DHCPv6 message	
DHCP Discover	Solicit	Sent from client. Tries to find any available DHCP server
DHCP Offer	Advertise	Sent from server to client in response to DHCP Discover or Solicit. Advertises that server is available for DHCP services
DHCP Request	Request	Sent from client to server. Requests IP address and other configuration parameters
	Renew	Sent from client to the server that originally provided the IP address. Asks to extend IP address lease
	Rebind	Sent from client to any available server if a past Renew message produced no answer. Asks to extend IP address lease
	Confirm	Sent from client to any available server. Asks to determine whether the allocated IP address is still valid on the link
DHCP ACK	Reply	Sent from server to client in response to multiple types of message. Content varies
DHCP Release	Release	Sent from client to the server that originally provided the IP address. Asks to cancel the IP address lease
DHCP Decline	Decline	Sent from client to server. Client has detected that the IP address assigned by the server is already in use
	Reconfigure	Sent from server to client. Server has new configuration parameters and the client is asked to initiate a Renew or Information-Request
DHCP Inform	Information-Request	Sent from client to server. Requests configuration parameters without any IP address
	Relay-Forward	Sent from relay agent to server or another relay agent. Content is another encapsulated message
	Relay-Reply	Sent from server to relay agent. Content is another encapsulated message
DHCP NAK		Sent from server to client. Client has incorrect parameters for the link or its lease has expired

PAM (Pluggable Authentication Modules) is an abstraction layer that allows applications to use authentication methods while being implementation-agnostic.

`/etc/pam.d/service`

PAM configuration for *service*

`/etc/pam.conf`

PAM configuration for all services. Obsolete

`ldd /usr/sbin/service | grep libpam`

Check if *service* is enabled to use PAM

<code>/etc/pam.d/service</code>		
auth	requisite	pam_securetty.so
auth	required	pam_nologin.so
auth	required	pam_env.so
auth	required	pam_unix.so nullok
account	required	pam_unix.so
session	required	pam_unix.so
session	optional	pam_lastlog.so
password	required	pam_unix.so obscure min=8 max=32
auth	required	pam_userdb.so db=/etc/vsftpd/logins
account	required	pam_userdb.so db=/etc/vsftpd/logins

type	auth	Authentication module to verify user identity and group membership
	account	Authorization module to determine user's right to access a resource (other than their identity)
	password	Module to update a user's authentication credentials
	session	Module (run at end and beginning of a user session) to set up the user environment
control	optional	Module is not critical to the success or failure of <i>service</i>
	sufficient	If this module succeeds, and no previous module has failed, module stack processing ends successfully. If this module fails, it is non-fatal and processing of the stack continues
	required	If this module fails, processing of the stack continues until the end, and <i>service</i> fails
	requisite	If this module fails, <i>service</i> fails and control returns to the application that invoked <i>service</i>
	include	Include modules from another PAM service file
module	PAM module and its arguments, e.g.:	
	pam_unix.so	Standard UNIX authentication module via <code>/etc/passwd</code> and <code>/etc/shadow</code>
	pam_nis.so	Module for authentication via NIS
	pam_ldap.so	Module for authentication via LDAP
	pam_radius_auth.so	Module for authentication via RADIUS
	pam_fshadow.so	Module for authentication against an alternative shadow passwords file
	pam_cracklib.so	Module for password strength policies (e.g. length, case, max number of retries)
	pam_limits.so	Module for system policies and system resource usage limits
	pam_listfile.so	Module to deny or allow the service based on an arbitrary text file
	pam_userdb.so	Module for authentication via a database. The option <code>db=/path</code> specifies the database location, which is a directory containing text files formatted as such: <i>username</i> <i>password</i> The user password is encrypted if the option <code>crypt=crypt</code> is used

LDAP (Lightweight Directory Access Protocol) is a simplified version of the X.500 standard and uses TCP port 389. LDAP allows organizing hierarchically a database of entries, each one of which is identified by a unique **DN (Distinguished Name)**. Each DN has a set of **attributes**, and each attribute has a **value**; an attribute may appear multiple times. Special attributes called **objectClass** define which attributes are allowed and which are required, and determine the **schema** of the LDAP.

dn: cn=John Doe,ou=IT Dept,dc=example,dc=org		Distinguished Name
Examples of LDAP attributes		
Attribute	Attribute with value	Meaning
cn	cn: John Doe	Common Name
dc	dc=example,dc=org	Domain Component
givenName	givenName: John	First name
sn	sn: Doe	Surname
mail	mail: jdoe@example.org	Email address
telephoneNumber	telephoneNumber: +1 555 1234 567	Telephone number
uid	uid: jdoe	User ID
c	c: US	Country code
l	l: San Francisco	Locality
st	st: California	State or province
street	street: 42, Penguin Road	Street
o	o: The Example Foundation	Organization
ou	ou: IT Dept	Organizational Unit
manager	manager: cn=Kim Green,ou=RD,dc=example,dc=org	Manager

LDIF (LDAP Data Interchange Format) is a plaintext data format for representing LDAP content and changes. The following LDIF file will change the email address of user "jdoe", add a picture, and delete the description attribute for the entry:

```
dn: cn=John Doe,dc=example,dc=org
changetype: modify
replace: mail
mail: johndoe@example.org
-
add: jpegPhoto
jpegPhoto:< file://tmp/jdoe.jpg
-
delete: description
-
```

<code>ldapsearch</code>	Query an LDAP server and return the output in LDIF
<code>-b base</code>	Start searching from <i>base</i>
<code>-z n</code>	Retrieve at maximum <i>n</i> entries as result
<code>-LLL</code>	Terse output. Outputs the result in LDIFv1, does not print comments, and omits the LDIF version number
<code>filter</code>	Search filter. If not specified, uses the default filter (<code>objectClass=*</code>)
<code>attributes</code>	Attributes to return. If not specified, returns all attributes
<code>ldapmodify</code>	Modify an LDAP entry
<code>ldapadd</code>	Add an LDAP entry
<code>ldapmodify -a</code>	
<code>ldapdelete</code>	Delete an LDAP entry
<code>-f file.ldif</code>	Modify, add, or delete an entry according to the LDIF file specified
<code>ldappasswd</code>	Change the password of an LDAP entry
<code>-s password</code>	Set the new password as <i>password</i>
<code>-S</code>	Prompt for the new password

In addition to the command-specific arguments, all LDAP commands above accept the following generic arguments:

<code>-H ldap://srv</code>	Connect to the specified LDAP server
<code>-H ldapi://</code>	Connect to the localhost LDAP server using IPC instead of a network socket
<code>-D binddn</code>	Bind (authenticate) to the LDAP server as the specified DN
<code>-w password</code>	Authenticate with the specified <i>password</i>
<code>-W</code>	Prompt for authentication
<code>-x</code>	Use simple authentication instead of SASL
<code>-v</code>	Use verbose mode for output

<code>ldapsearch -H ldap://ldap.example.org \</code> <code>-s base -b "ou=people,dc=example,dc=com" "(sn=Doe)" \</code> <code>cn sn telephoneNumber</code>	Query a LDAP server for entries in the OU "people" whose surname is "Doe"; print common name, surname, and telephone number of the entries found
<code>ldapmodify -b -r -f file.ldif</code>	Modify an entry according to the LDIF file specified
<code>ldapadd -h ldap.example.org \</code> <code>-D "cn=Admin,dc=example,dc=org" -W -f file.ldif</code>	Authenticating as "Admin", add an entry by adding the content of the specified LDIF file to the directory
<code>ldapdelete -h ldap.example.org \</code> <code>-D "cn=Admin,dc=example,dc=org" -W \</code> <code>"uid=jdoe,dc=example,dc=org"</code>	Authenticating as "Admin", delete the user "jdoe"
<code>ldappasswd -h ldap.example.org \</code> <code>-D "cn=Admin,dc=example,dc=org" -W -x \</code> <code>-S "uid=jdoe,ou=IT Dept,dc=example,dc=org"</code>	Authenticating as "Admin" on example.org, change the password of user "jdoe" in the OU "IT Dept"

OpenLDAP is an open source implementation of LDAP, and was initially developed together with the LDAP protocol. Its related service is `slapd`, the Standalone LDAP daemon. SSSD can be configured to provide access to OpenLDAP (or any other LDAP server) as an authentication and identity provider.

<code>/var/lib/ldap/</code>	Files constituting the OpenLDAP database
<code>/etc/openldap/slapd.conf</code> <code>/usr/local/etc/openldap/slapd.conf</code>	OpenLDAP configuration file (deprecated)
<code>/usr/local/etc/openldap/slapd.d/</code> (v2.3 and later)	Directory containing the LDIF database that stores the OpenLDAP configuration. These LDIF files must not be edited by hand
<code>slaptest -u</code>	Verify that the OpenLDAP configuration is correct
<code>slapcat -l file.ldif</code>	Dump the contents of an OpenLDAP database to an LDIF file
<code>slapadd -l file.ldif</code>	Import an OpenLDAP database from an LDIF file
<code>slapindex</code>	Regenerate OpenLDAP's database indexes
<code>slapcat -b cn=config</code> <code>ldapsearch -Y EXTERNAL -H ldapi:/// -b cn=config</code>	Show the OpenLDAP configuration
<code>yum install openldap openldap-clients authconfig \</code> <code>sssd nss-pam-ldapd authconfig-gtk</code> (RHEL 7)	Install the OpenLDAP client
<code>authconfig --enableldap --enableldapauth \</code> <code>--ldapserver=ldap://ldapserver \</code> <code>--ldapbasedn="dc=example,dc=org" \</code> <code>--enablenesssd --update</code> (RHEL 7)	Set up the LDAP client to connect to a <i>ldapserver</i> . This will update the configuration files <code>/etc/sss/sss.conf</code> and <code>/etc/openldap/ldap.conf</code>
<code>authselect select sssd --force</code> (RHEL 8)	Set up LDAP client authentication via <code>sss</code>
<code>authconfig-gtk</code> <code>system-config-authentication</code>	OpenLDAP configuration GUI

389 Directory Server is an enterprise-class open source LDAP server. It is derived from OpenLDAP and is part of the Fedora Project. A commercial version is also available with the name **Red Hat Directory Server**.

It features TLSv1 encryption, SASL, synchronization with MS Windows Active Directory, and a web console (on port 9090). It also includes Lib389, a Python base library that can be used to manage, test, and perform all operations on a 389 DS instance.

<code>/etc/dirsrv/slapd-<i>instancename</i>/dse.ldif</code>	Instance configuration (<code>cn=config</code> entry)
<code>/var/lib/dirsrv/slapd-<i>instancename</i>/</code>	Directory containing the database and other data relative to an instance

<code>dsctl</code>	Start, stop, display status, backup, and generally manage a local instance
<code>dsconf</code>	Configure a local or remote instance
<code>dsidm</code>	Manage backend data (users, groups, permissions)

SELinux (Security-Enhanced Linux) is a Linux kernel security module that provides a mechanism for supporting access control security policies.

SELinux implements a Mandatory Access Control framework that allows the definition of fine-grained permissions for how **subjects** (i.e. processes) access **objects** (i.e. other processes, files, devices, ports, sockets); this improves security with respect to the traditional Discretionary Access Control, which defines accesses based on users and groups. Processes, files, and users have a **security context** structured as *user:role:type:level* e.g. `unconfined_u:object_r:user_home_t:s0`. The third field defines a *type* for files or a *domain* for processes. The security context of a file is stored in its extended attributes. The decisions SELinux takes about allowing or disallowing access are stored in the **AVC (Access Vector Cache)**.

SELinux creates a pseudo filesystem (SELinuxfs) containing commands used by the kernel for its operations; this filesystem is usually mounted on `/selinux/` or `/sys/fs/selinux/`.

```
setenforce 0
echo 0 > /selinux/enforce

setenforce 1
echo 1 > /selinux/enforce

getenforce
cat /selinux/enforce
```

Enter permissive mode (SELinux must be enabled)

Enter enforcing mode (SELinux must be enabled)

Display current mode

```
sestatus -v
```

Show SELinux mode, SELinuxfs mount point, etc.

```
/etc/selinux/config
/etc/sysconfig/selinux (symlink)
```

Configuration file indicating the permanent SELinux state:

mode	SELINUX=	enforcing	SELinux fully enforces security policies
		permissive	SELinux does not enforce security policies, but logs all violations
		disabled	SELinux security policies are disabled
policy	SELINUXTYPE=	targeted	SELinux protects targeted daemons
		strict	(up to RHEL 6) SELinux fully protects the system
		minimum	(RHEL 7 and later) SELinux only protects selected processes
		mls	(RHEL 7 and later) Multi Level Security protection

```
ls -Z
```

List files and their security context

```
ps -eZ
```

List processes and their security context

```
cp --preserve=context file file2
```

Copy a file and its security context.
By default, the `cp` command creates a new SELinux file context

```
tar --selinux otherargs
star -xattr -H=exustar otherargs
```

Create or extract archives that retain the security context of the original files

semanage

Manage SELinux policies

semanage fcontext -l

List files and their assigned SELinux labels

semanage fcontext -a -t *label file*

Assign the SELinux *label* to *file*. Afterwards, it is necessary to apply the label via `restorecon -f file`

semanage fcontext -a -t httpd_sys_content_t \
"/var/www2/html2(/.*)?"

Allow a local webserver to serve content stored in the directory `/var/www2/html2`

semanage login -l

List mappings between users and SELinux users

semanage port -l

List port numbers and their assigned SELinux type definitions

semanage port -a -t *portlabel* -p tcp *n*

Assign the SELinux *portlabel* to TCP port *n*

semanage port -a -t http_port_t -p tcp 8888

Allow a local webserver to serve content on port 8888

semanage port -d -t http_port_t -p tcp 8888

Remove the binding of `http_port_t` port label to TCP 8888

semanage port -m -t http_cache_port_t -p tcp 8888

Modify the port label bound to TCP 8888

semanage permissive -a auditd_t

Add `auditd_t` to the list of permissive types/domains. In this case, SELinux allows the `auditd` daemon all access while logging its AVC violations

semanage permissive -d auditd_t

Delete `auditd_t` from the list of permissive types/domains

semanage permissive -l

List all permissive types/domains

chcon *context file*

Change the security context of *file* to the specified *context*. Changes made via `chcon` are not persistent across filesystem relabels or the execution of `restorecon`; for persistent changes, use `semanage fcontext` followed by `restorecon`

chcon --reference=*file0 file*

Change the security context of *file* to be the same as *file0*

restorecon -f *file*

Restore the security context of *file* to the system default

matchpathcon *path*

Get the default security context for *path* from the file contexts configuration

getsebool *boolean*

Get the value of a SELinux boolean

setsebool -P *boolean=value*

Set the value of a SELinux boolean, writing it to the policy file on disk so that it will be persistent after reboot

AVC denials (i.e. SELinux violation events) have the format `type=AVC msg=audit(timestamp:id): avc: denied (...)`. They are logged to `/var/log/audit/audit.log` if the `auditd` daemon is running (default), or to `/var/log/messages` if the `rsyslogd` daemon is running. AVC denials can also be read via `dmesg`.

<code>sepolicy</code>	Inspect a SELinux policy
<code>sepolicy manpage -a -p /usr/local/man/man8 && mandb</code>	Generate all SELinux policy manpages
<code>seinfo</code>	Query the components of a SELinux policy
<code>sealert -a logfile</code>	Analyze a SELinux logfile and display verbosely SELinux policy violations, suggesting the solution (set a boolean, create and install the appropriate module, etc.)
<code>audit2why -d</code>	Read AVC violations from the output of <code>dmesg</code>
<code>grep timestamp:id logfile audit2why</code>	Diagnostic a specific AVC denial event entry (identified by a <i>timestamp</i> and an <i>id</i>) from a SELinux <i>logfile</i>
<code>audit2allow -i inputfile -M module</code>	Generate a loadable <i>module</i> containing the appropriate SELinux policy from a denied operation stored in <i>inputfile</i>
<code>fgrep "component" /var/log/audit/audit.log \</code> <code>audit2allow -m module.pp -o module.te</code>	Search the AVC denial log for occurrences of <i>component</i> , and generate the appropriate module <i>.pp</i> (Policy Package) appending the output to a <i>.te</i> file (Type Enforcement policy)
<code>ausearch -a id</code>	Query the SELinux log for event <i>id</i>
<code>ausearch -m AVC,USER_AVC,SELINUX_ERR -ts today</code>	Query the SELinux log for events matching the given message types happened today
<code>ausearch -c 'exe' --raw audit2allow -M module</code>	Generate a loadable module to allow access on an executable <i>exe</i> which caused an AVC violation. This module must then be installed via <code>semodule -i module.pp</code>
<code>checkmodule -M -m module.te -o module.mod</code>	Check and compile a SELinux non-base policy module (<i>.te</i>) into a binary representation (<i>.mod</i>)
<code>semodule_package -m module.mod -o module.pp</code>	Package a binary policy module (<i>.mod</i>) into a SELinux policy module package (<i>.pp</i>)
<code>semodule -l</code>	List installed SELinux policy modules
<code>semodule -X n -i module.pp</code>	Install a SELinux policy module at priority <i>n</i> . Installed modules persist after reboot
<code>semodule -X n -r module</code>	Remove a SELinux policy module at priority <i>n</i> . Modules must be removed at the same priority at which they were installed

AVC message format	
type=AVC	The entry is an Access Vector Cache message
msg=audit(1689243601.955:409):	Timestamp in Unix epoch time (1689243601.955) and unique identifier (409) of the message
avc: denied	The action was denied by SELinux
{ name_connect }	Type of access that was denied; in this case, a <code>name_connect</code> operation, which typically involves resolving a hostname to an IP address and establishing a network connection
for pid=1098	PID (1098) of the process that triggered the SELinux denial (<code>php-fpm</code>)
comm="php-fpm"	Name of the process that triggered the SELinux denial
dest=3307	Destination port (3307) to which the process was attempting to connect
scontext=system_u:system_r:httpd_t:s0	SELinux security context (<code>httpd_t</code>) of the source process; in this case, the process was likely part of the Apache HTTP Server <code>httpd</code>
tcontext=system_u:object_r:unreserved_port_t:s0	SELinux security context (<code>unreserved_port_t</code>) of the destination port; in this case, this indicates that port 3307 is not associated with a specific service and is considered an unreserved port
tclass=tcp_socket	Type of object (<code>tcp_socket</code>) that the process was attempting to access
permissive=1	SELinux is operating in permissive mode, which means that denials are logged but not enforced

`auditd` is the Linux Auditing System daemon, developed and maintained by Red Hat. It is used by SELinux to log events.

<code>auditctl</code>	Control and query the kernel audit system
<code>auditctl -a exit,always -S open -F path=<i>file</i></code>	Audit processes trying to open <i>file</i>
<code>auditctl -a exit,always -S open -F uid=<i>uid</i></code>	Audit files opened by the user with UID <i>uid</i>
<code>auditctl -w <i>file</i> -p wa</code> <code>auditctl -a exit,always -F path=<i>file</i> -F perm=wa</code>	Audit <i>file</i> for changes
<code>ausearch -f <i>file</i></code>	Search the audited events matching <i>file</i>
<code>aureport</code>	Produce a summary report about audited events

Tor (The Onion Routing) is an overlay network which provides anonymity over the Internet, masking the user's IP address.

It is based on open source software and on the concept of **onion routing**, i.e. a communication technique where messages are encapsulated in layers of encryption, with each layer being encrypted or decrypted as the message hops from a router to another.

Tor is an important component of privacy-savvy Linux distros such as Tails and Whonix.

The Tor client (**onion proxy**) establishes an encrypted **circuit** through three Tor relays (chosen randomly from a directory of publicly available Tor relays aka **onion routers**): an **entry guard** aka **guard relay**, a **middle relay**, and an **exit relay**. The exit relay transmits the communication to the intended destination, which sees the exit relay's IP address as the source of the connection. Only the traffic between the exit relay and the destination is non-encrypted.

The Tor client can also use a **bridge relay** in lieu of a guard relay. Bridge relays function similarly to Tor relays but are not listed publicly, so they cannot be blacklisted by a censoring government. Bridge relays help in concealing the usage of Tor but may slow down the connection.

Once the Tor service is started, applications (e.g. web browsers) that need to use it must have their SOCKS (V5) proxy manually configured to localhost on port 9050. Any active content (JavaScript, ActiveX, etc.) must be disabled in the web browser, as it connects directly to the destination server and will leak the client's real IP address. Note that Tor does not route UDP traffic.

Onion Services are services that can be accessed only through Tor. They have an onion address `string.onion` where `string` is the encoded identity public key of the Onion Service.

In the setup phase, an Onion Service establishes long-term circuits with three Tor relays, which will act as exclusive introduction points to the service. Then the Onion Service creates a descriptor containing the list of its introduction points and their authentication keys, signs it with the Onion Service's identity private key, and uploads it to a distributed hash table within the Tor network.

A client wanting to access the Onion Service searches for its descriptor in the distributed hash table, verifies its signature using the public key encoded in the onion address, and contacts the service via one of the introduction points. Then the client establishes the usual three-node Tor circuit, with the third relay to be used as a rendezvous point. The Onion Service connects to the rendezvous point through another three-node Tor circuit, and can start communicating with the client.

<code>tor --version</code>	Display the version of Tor
<code>systemctl start tor.service</code>	Start the Tor service
<code>torsocks</code>	Allows to use most applications through Tor. It safely handles DNS requests and rejects non-TCP traffic from the torified application
<code>tor-router</code>	Allows to use Tor for all network traffic, including DNS requests

Kickstart is a method to perform automatic installation and configuration of RHEL machines.

This can be done by specifying `inst.ks=hd:/dev/sda:/root/path/ksfile` either as a boot option, or an option to the kernel command in GRUB 2.

<code>/root/anaconda-ks.cfg</code>	Kickstart file describing the current system. This file is automatically generated during the installation
<code>system-config-kickstart</code>	GUI tool to create a Kickstart file
<code>ksvalidator ksfile</code>	Check the validity of a Kickstart file
<code>ksverdiff -f RHEL6 -t RHEL7</code>	Show the differences in the Kickstart syntax between RHEL 6 and RHEL 7

Red Hat Satellite is a system management software platform that allows provisioning and configuration of RHEL machines. Repository content is provided via Red Hat Subscription Management (RHSM). Satellite 5 was based on **Spacewalk**, an open source system management software for Linux machines. Satellite 6 is a complete overhaul of it and is based on Foreman as its upstream open source project.

Foreman is an open source lifecycle management tool able to provision servers and manage their configuration; it supports Kickstart, Puppet, Ansible, Chef, and Salt.

Katello is a Foreman plugin that handles Red Hat repository management (via the **Pulp** service) and subscription management (via the **Candlepin** service). Starting with Satellite 6.9, the Katello agent is deprecated and is going to be removed from future releases; workloads will use Remote Execution only. All these components above need a PostgreSQL database, except Pulp which needs a MongoDB database.

As a separate component, **Capsule** servers act as proxies for many of the main Satellite functions e.g. repository storage. A Capsule is also integrated in each Satellite server.

The equivalent of a Capsule in Foreman is a **Smart Proxy**.

<code>subscription-manager register --username user --password psw</code>	Register a system to the RHSM portal through the account identified by <i>user</i>
<code>subscription-manager attach --auto</code>	Attach a RHSM subscription to a registered system
<code>subscription-manager unregister</code>	Unregister a system from RHSM
<code>subscription-manager clean</code>	Detach a RHSM subscription from a system
<code>foreman-maintain service list</code>	List all Satellite services
<code>foreman-maintain service status</code>	Display status or start, stop, restart all Satellite services.
<code>foreman-maintain service start</code>	Performed via <code>systemctl</code>
<code>foreman-maintain service stop</code>	
<code>foreman-maintain service restart</code>	
<code>foreman-maintain backup</code>	Make a backup of Satellite
<code>foreman-rake command:option</code>	Perform various administrative tasks
<code>hammer</code>	CLI tool for Foreman
<code>pulp-admin-client</code>	Tool to administer the Pulp server
<code>virt-who</code>	Agent for scanning and reporting virtual guest IDs and hypervisors to a Satellite server
<code>foreman-debug</code>	Collect Satellite configuration, log, and backend data for debug purposes

Virtualization software technologies (KVM, VMware, Xen, User-mode Linux, etc.) allow running a virtual instance of a system (i.e. a Virtual Machine) in a layer abstracted from the actual hardware. A **hypervisor** (either bare-metal aka type-1 or software aka hosted aka type-2) installed on the host machine allows running multiple guest OSes with different kernels and their applications; these OSes coexist separately as they were on dedicated physical machines. In **full virtualization** the hardware is fully simulated. In **paravirtualization**, hardware is not simulated; guest applications are executed in their own isolated domains as if they were running on separate systems, but need to be specifically modified to run in that environment.

KVM (Kernel-based Virtual Machine) is a virtualization infrastructure for the Linux kernel that allows it to function as a hypervisor. It was introduced in version 2.6.20 of the Linux kernel.

Red Hat Virtualization, formerly known as Red Hat Enterprise Virtualization (RHEV), is based on KVM.

<code>/etc/libvirt/qemu/</code>	Directory containing the XML files that define VMs properties. <code>libvirtd</code> must be restarted after modifying an XML file
<code>/var/lib/libvirt/</code>	Directory containing files related to the VMs
<code>virt-manager</code>	KVM GUI
<code>virt-install --prompt</code>	Interactive command-line program to create a VM
<code>virt-install -n <i>vmname</i> -r 2048 \</code> <code>--disk path=/var/lib/libvirt/images/<i>vmname</i>.img \</code> <code>-l /root/vmstuff/inst/ \</code> <code>-x "ks=/root/vmstuff/kickstart.cfg"</code>	Create a VM with 2 Gb of RAM, specifying path of virtual disk, location of installation files, and (as extra argument) the Kickstart configuration to use
<code>virt-clone --prompt</code>	Interactive command-line program to clone a VM. A VM must be shut off or paused before it can be cloned
<code>virt-clone -o <i>vmname</i> -n <i>vmclonename</i></code>	Clone a VM
<code>virsh</code>	Interface for VM management
<code>virsh list --all</code>	List all VMs present on the system
<code>virsh start <i>vmname</i></code>	Start a VM
<code>virsh destroy <i>vmname</i></code>	Brutally shut down a VM
<code>virsh shutdown <i>vmname</i></code>	Gracefully shut down a VM
<code>virsh autostart <i>vmname</i></code>	Set a VM to be automatically started when the system boots. Done by symlinking the VM to <code>/etc/libvirt/qemu/autostart/</code>
<code>virsh autostart --disable <i>vmname</i></code>	Disable the autostart of a VM at system boot
<code>virsh edit <i>vmname</i></code>	Edit the XML file defining a VM's properties
<code>virt-what</code>	Detect whether the current machine is a VM

QEMU (The Quick Emulator) is an open source machine emulator and virtualizer that uses dynamic binary translation to emulate a computer's processor. It can run a whole virtual machine as well as run applications compiled for another CPU architecture.

<code>qemu-system-alpha</code>	Run a virtual machine with alpha processor
<code>qemu-system-arm</code>	Run a virtual machine with ARM processor
<code>qemu-system-ppc</code>	Run a virtual machine with HPPA architecture
<code>qemu-system-i386</code>	Run a virtual machine with Intel 32 bit x86 CPU
<code>qemu-system-x86_64</code>	Run a virtual machine with Intel 64 bit x86 CPU
<code>qemu-system-lm32</code>	Run a virtual machine with LatticeMico32 CPU
<code>qemu-system-m68k</code>	Run a virtual NeXT Cube, Macintosh Quadra 800, MCF5208EB, or Arnewsh 5206
<code>qemu-system-microblazeel</code>	Run a virtual machine with MicroBlaze RISC CPU
<code>qemu-system-mips</code>	Run a virtual machine with a MIPS Malta system
<code>qemu-system-nios2</code>	Run a virtual machine with Nios2 CPU
<code>qemu-system-or1k</code>	Run a virtual machine with OpenRISC architecture
<code>qemu-system-ppc</code>	Run a virtual PowerPC system
<code>qemu-system-riscv64</code>	Run a virtual machine with RISC-V architecture
<code>qemu-system-sh4</code>	Run a virtual machine with SH-4 RISC CPU
<code>qemu-system-sparc</code>	Run a virtual machine with Sun SPARC architecture
<code>qemu-system-s390x</code>	Run a virtual s390x zSystems mainframe
<code>qemu-system-tricore</code>	Run a virtual machine with TriCore architecture
<code>qemu-system-unicore32</code>	Run a virtual machine with Unicore architecture
<code>qemu-system-xtensa</code>	Run a virtual machine with Tensilica Xtensa processor
<code>qemu-alpha</code>	Execute an alpha binary
<code>qemu-arm</code>	Execute an ARM binary
<code>qemu-armeb</code>	
<code>qemu-m68k</code>	Execute a Motorola 68000 binary
<code>qemu-i386</code>	Execute an Intel 32 bit x86 binary
<code>qemu-x86_64</code>	Execute an Intel 64 bit x86 binary
<code>qemu-microblaze</code>	Execute a MicroBlaze binary
<code>qemu-mips</code>	Execute a 32 bit big endian MIPS binary (MIPS O32 ABI)
<code>qemu-mipsel</code>	Execute a 32 bit little endian MIPS binary (MIPS O32 ABI)
<code>qemu-mips64</code>	Execute a 64 bit big endian MIPS binary (MIPS N64 ABI)
<code>qemu-mips64el</code>	Execute a 64 bit little endian MIPS binary (MIPS N64 ABI)
<code>qemu-mipsn32</code>	Execute a 32 bit big endian MIPS binary (MIPS N32 ABI)
<code>qemu-mipsn32el</code>	Execute a 32 bit little endian MIPS binary (MIPS N32 ABI)
<code>qemu-ppc</code>	Execute a PowerPC binary
<code>qemu-ppc64</code>	
<code>qemu-sh4</code>	Execute an SH-4 binary
<code>qemu-sh4eb</code>	
<code>qemu-sparc</code>	Execute a SPARC32 binary
<code>qemu-sparc32plus</code>	Execute a SPARC32 or SPARC32PLUS binary
<code>qemu-sparc64</code>	Execute a SPARC64 or SPARC32PLUS binary

Multipass is a tool to easily generate and manage cloud-style lightweight Ubuntu VMs.

<code>multipass</code>	Multipass CLI
<code>multipass launch</code>	Create a VM with default settings and a random name, using the latest Ubuntu LTS release as the base image
<code>multipass launch ubuntu:24.04 \</code> <code>--name <i>vmname</i> --mem 4G --disk 40G</code>	Create a VM using Ubuntu 24.04 as the base image, with 4 Gb RAM and 40 Gb disk space
<code>multipass start <i>vmname</i></code>	Start a VM
<code>multipass stop <i>vmname</i></code>	Stop a VM
<code>multipass info <i>vmname</i></code>	Display information about a VM
<code>multipass shell <i>vmname</i></code>	Login to a VM and use their shell prompt
<code>multipass exec <i>vmname</i> -- /bin/bash</code>	Launch a Bash shell on a VM
<code>multipass transfer /localpath/file \</code> <code><i>vmname</i>:/vmpath</code>	Transfer a file from the local filesystem to a VM's disk
<code>multipass delete <i>vmname</i></code>	Delete a VM

In **containerization** aka **OS-level virtualization** (Docker, rkt, OpenVZ, Virtuozzo, etc.) the kernel allows the existence of multiple isolated user space instances, called **containers**. A container is a sandboxed software application packaged with all its dependencies and its configuration so that it is able to run in any environment. There is no hypervisor; all containers are run by the container runtime engine, which is placed on top of the OS.

Containerization is more lightweight and efficient than virtualization, because programs in OS-level virtual partitions can use the OS's normal system call interface, with no need for emulation. However, it is also less flexible as it can only run guest OSes that share the same kernel version. So a running container always uses the kernel of the host operating system. It is possible for container images for different Linux distributions to coexist because these images do not include their own kernels; instead, they contain only the userland components of the respective distributions e.g. system libraries, package managers, and other utilities specific to that distribution.

The original userspace interface for the Linux kernel containment feature is **LXC (Linux Containers)**.

Docker is a PaaS that implements containerization.

Docker uses a **Dockerfile** as a blueprint to build an **image**, which is a read-only file containing the source code, libraries, and dependencies of an application. Images are pushed (i.e. stored) or pulled (i.e. retrieved) from a **container registry** aka **Docker registry**. Then, the image is run, resulting in a **container** (in execution).

Docker registries are usually cloud-based; the default registry is Docker Hub, the primary and largest library and community for Docker images. A **Docker repository** is a collection of related Docker images within a registry, whereas tags are used to differentiate between multiple versions or variants of an image.

Images are composed of multiple layers, and common layers are shared between images; this optimizes disk space and network bandwidth when pushing or pulling large images.

All changes to the filesystem of a Docker container (files created, modified, or deleted) are isolated to that container and lost when the container is removed. To ensure persistence of files, a **volume** can be used to connect from a specific filesystem path of the container to the filesystem on the host machine (`/var/lib/docker/volumes/` by default). Another possibility is to use a **bind mount**, which allows to directly mount a host directory into the container; this is used mostly for temporary storage.

Creating a Docker **network** allows to have multi-container applications, as containers are able to communicate only if they are in the same network.

Docker Compose is a tool to define and share multi-container applications. It uses a YAML file to define all services so that they can be operated with a single command.

<code>docker</code>	Docker CLI
<code>docker info</code>	Display information about the Docker installation
<code>docker pull image</code>	Pull an image from the previously specified registry
<code>docker image ls</code>	List the created images
<code>docker image history image</code>	Show the layers composing an image
<code>docker ps</code>	List all running containers
<code>docker ps -a</code>	List all containers that ran and exited successfully
<code>docker run image:v1</code>	Run an image as a container
<code>docker build -t image:v1 /path</code>	Create an image from the Dockerfile, using <i>path</i> (where the Dockerfile is located) as the build context
<code>docker push image:v1</code>	Push an image to the previously specified registry
<code>docker scan image</code>	Scan an image for security vulnerabilities (via Snyk)
<code>docker tag image:v1 image2:v1</code>	Create a new tag for the same image. Both images will have the same image ID
<code>docker inspect containerID</code>	Inspect a container
<code>docker start containerID</code>	Start a container
<code>docker stop containerID</code>	Stop a container
<code>docker stop \$(docker ps -q)</code>	Stop all running containers
<code>docker restart containerID</code>	Restart a container
<code>docker logs containerID</code>	Show the logs for a container
<code>docker rm containerID</code>	Remove a container. A container must be in a stopped state to be removed
<code>docker volume create volname</code>	Create a named volume
<code>docker volume inspect volname</code>	Show info about a named volume, including its mountpoint on the host machine
<code>docker network create app</code>	Create a network

Dockerfile example	
<code>FROM ubuntu:18.04</code>	Create a layer from the "ubuntu:18.04" Docker image
<code>COPY . /app</code>	Add files from the current directory on the Docker client
<code>RUN make /app</code>	Run the "make" command to build the application
<code>CMD python /app/app.py</code>	Run a command inside the container

<code>hadolint</code>	Dockerfile linter tool. Checks a Dockerfile for errors, security vulnerabilities, and performance problems
<code>trivy</code>	Security tool to find vulnerabilities and misconfigurations in container images, Kubernetes clusters, code repositories, binary artifacts, etc.

Container orchestration helps in the provisioning and deployment of containers, provides scaling and load balancing, ensures redundancy and availability, and allows performing scheduling and health checks.

Kubernetes (aka **K8s**) is the most used container orchestration platform. It is open source software.

A **Kubernetes cluster** is composed of the **Control Plane**, which manages the cluster, and of several nodes, which are the workers that host the running applications. A Kubernetes cluster in a production environment should have a minimum of three nodes.


Nodes can be VMs or physical machines. Nodes run pods, which consist of the containerized applications. Each node has also a **Kubelet**, which is an agent that manages and monitors the node and communicates with the Control Plane, and a container runtime tool (e.g. Docker or containerd) for pulling the container image from a registry, unpacking the container, and running the application.

A **Deployment** instructs Kubernetes how to create instances of an application. Once the application has been deployed, Kubernetes continuously monitor it to ensure that any node that goes down is replaced with another node, through a self-healing mechanism.

A **pod** is the atomic unit of deployment and consists of one or more containers running in a shared context, sharing storage and network resources, and having a common specification for how to run the containers. Usually a pod hosts only one container. However, in case of tightly coupled applications e.g. a WordPress web application with its MySQL database, it makes more sense to host the web application (**app container**) and the database (**sidecar container**) in the same pod. In other cases, it is useful to host in the same pod additional containers that run and complete in order (**init containers**) before the main app container runs.

Each pod has a unique IP address, so different containerized application within the same pod communicate between them via localhost on different ports.

Pods are not created directly but are generated via workload resources e.g. Deployments. A pod remains on a node until the pod finishes execution, is deleted, is removed for lack of resources, or the node fails.



Volumes allow containers in pods to access data through a filesystem, and can be persistent or ephemeral i.e. linked to the lifetime of a specific pod. Kubernetes supports many kinds of volumes: emptyDir (pod-based volume), hostPath (a directory on the host node's filesystem mounted on the pod), local (dedicated local storage device on a node, with pods using it always scheduled to that same node), image (Oracle Cloud Infrastructure object storage), iSCSI volumes, fibre channel, NFS, etc.

A **ConfigMap** is an object that stores data in key-value pairs, and is mostly used to separate configuration data from code. A pod can access a ConfigMap as an environment variable, as command line arguments, or as a volume.

A **Secret** is similar to a ConfigMap, but is used to hold confidential data; e.g. this data is not printed when using the `kubectl describe` command.



A **Lease** provides a mechanism to lock shared resources, coordinate activity between members of a set, communicate node heartbeats, and implement leader election in HA configurations.

Namespaces allow isolating groups of resources within a cluster. The four predefined Kubernetes namespaces are: default (default namespace ready for user objects), kube-node-lease (holds Lease objects), kube-public (publicly readable by all clients), and kube-system (holds objects created by the Kubernetes system).

A **Service** defines a logical set of pods and a policy to access them. A Service also allow the pods' IP address to be exposed outside the cluster, and the applications to exchange network traffic.

Services match a set of pods using **labels** and **label selectors**. Labels are key/value pairs attached to objects to define their environment (production, test, etc.), version, or object type.

A **ReplicaSet** ensures that a specified number of identical redundant pods are running at any given time; new pods will be created and assigned to nodes with available resources so to increase the number of pods to the desired state. A ReplicaSet is usually managed automatically by a Deployment.



<code>kubectl</code>	Kubernetes CLI to interact with the Control Plane
<code>kubectl run nginx --image nginx</code>	Create a pod running a Nginx container
<code>kubectl create -f nginx.yaml</code>	Create an object based on a configuration template
<code>kubectl apply -f nginx/</code>	Apply all files in a directory
<code>kubectl exec podname -- command</code>	Run <i>command</i> on a pod
<code>kubectl exec -ti podname -- bash</code>	Run a Bash shell session on a pod
<code>kubectl get resources</code>	List the specified resources
<code>kubectl get nodes</code>	List the available nodes in the cluster
<code>kubectl get pods</code>	List pods
<code>kubectl describe resources</code>	Display detailed information about the specified resources
<code>kubectl cluster-info</code>	Display information about the cluster
<code>kubectl logs podname</code>	Display logs for a specific pod
<code>kubectl proxy</code>	Applications run inside an isolated Kubernetes network, which allows communications only between pods and services within the same cluster. This command (which must be run on a different terminal) creates a proxy that will forward communications outside the private network, and e.g. will allow a connection between the host terminal and the Kubernetes cluster
<code>kubeadm</code>	Tool for quickly bootstrapping a cluster. It does not include options for machine provisioning
<code>k9s</code>	Terminal-based UI for easy management of a Kubernetes cluster

How to deploy an app

1. `kubectl create deployment kubernetes-bootcamp \`
`--image=gcr.io/google-samples/kubernetes-bootcamp:v1`
2. `kubectl proxy`
3. `curl http://localhost:8001/version`
4. `export POD_NAME=$(kubectl get pods -o go-template --template \`
`'{{range .items}}{{.metadata.name}}\n'{{end}}')`
5. `curl http://localhost:8001/api/v1/namespaces/default/\`
`pods/$POD_NAME:8080/proxy/`

Create a Deployment. This schedules the application to run on a suitable node

Start the proxy (in another terminal)

Query the version through the proxied API

Get the pod's name

Access the pod through the proxied API

How to expose an app to the outside network through a service

1. `kubectl get services`
2. `kubectl expose deployment/kubernetes-bootcamp \`
`--type="NodePort" --port 8080`
3. `kubectl describe services/kubernetes-bootcamp`
4. `export NODE_PORT=$(kubectl get services/kubernetes-bootcamp \`
`-o go-template='{{(index .spec.ports 0).nodePort}}')`
5. `curl http://$(minikube ip):$NODE_PORT`

List the current Services

Expose a deployment to the outside network by creating a Service of type NodePort with a unique cluster IP address, internal port, and external IP address (equivalent to the IP address of the node)

Examine the service

Get the network port on which the node exposes the service

Access the service through the node's IP address and port

How to deploy multiple instances of an app

1. `kubectl expose deployment/kubernetes-bootcamp \`
`--type="LoadBalancer" --port 8080`
2. `kubectl get deployments`
3. `kubectl get rs`
4. `kubectl scale deployments/kubernetes-bootcamp --replicas=4`
5. `kubectl get pods -o wide`
6. `kubectl describe deployments/kubernetes-bootcamp`
7. `export NODE_PORT=$(kubectl get services/kubernetes-bootcamp \`
`-o go-template='{{(index .spec.ports 0).nodePort}}')`
8. `curl http://$(minikube ip):$NODE_PORT`

Create a Service of type LoadBalancer and expose it

List the current Deployment

List the ReplicaSet created by the Deployment

Scale the Deployment to 4 replicas

Verify that 4 pods are now active

Display the Deployment event log

Get the network port on which the node exposes the service

Access the service through the node's IP address and port. Verify that a different pod answers the request each time this command is run

minikube is a lightweight user-friendly small-scale version of Kubernetes. By default, it creates a single-node local Kubernetes cluster. It has the purpose of making it easy to learn, develop, and test Kubernetes.

<code>minikube</code>	minikube CLI
<code>minikube ip</code>	Display the IP address of the primary control plane
<code>minikube ip -n <i>node</i></code>	Display the IP address of <i>node</i>
<code>minikube service <i>service</i></code>	Start a tunnel exposing <i>service</i> , displaying the local port used on localhost

MicroK8s is a lightweight distribution of Kubernetes developed by Canonical.

<code>microk8s</code>	MicroK8s CLI
<code>microk8s status</code>	Display the status of MicroK8s
<code>microk8s enable <i>addon</i></code>	Enable an add-on
<code>microk8s disable <i>addon</i></code>	Disable an add-on
<code>microk8s kubectl <i>command</i></code>	Execute a Kubernetes command
<code>microk8s config</code>	Display the Kubernetes configuration file

Cloud computing is the on-demand delivery of computing services through the Internet. The cloud provider supplies these services according to different models:

IaaS (Infrastructure as a Service)	Provides virtual machines, storage, load balancing, networking, etc.
PaaS (Platform as a Service)	Provides OS, databases, development environments, web servers, etc.
SaaS (Software as a Service)	Provides access to ready-to-use application software
Serverless computing	Provides computing execution power, by provisioning computing resources (virtual machines, storage, and network) as needed, transparently to the user

In cloud-native development, a **service mesh** is a dedicated layer to make communications between microservices secure and reliable.

CI/CD (Continuous Integration / Continuous Delivery or Deployment) is a method to frequently deploy applications in production. It is based on the automation of new code build, test, merge, release to repository, and deployment.

OpenStack is an open standard cloud computing platform, usually deployed as a IaaS solution for public and private clouds. It has a modular architecture, composed of many elements:

- **Aodh**, a rule-based alarm service;
- **Barbican**, a REST API for management of keys and secrets;
- **Ceilometer**, for telemetry;
- **Cinder**, a block storage service;
- **Designate**, a REST API for DNS management;
- **Glance**, a service to upload and distribute images;
- **Heat**, an orchestration service;
- **Horizon**, a web-based dashboard;
- **Ironic**, for provisioning bare metal servers;
- **Keystone**, for identity and authentication;
- **Magnum**, for container orchestration;
- **Manila**, a shared file system;
- **Mistral**, to manage workflows;
- **Neutron**, to manage networking;
- **Nova**, for provisioning of computing instances (either VMs or bare metal servers);
- **Sahara**, to provision Hadoop clusters;
- **Searchlight**, an Elasticsearch-based search tool for OpenStack cloud services;
- **Swift**, a distributed object store;
- **Trove**, a relational and non-relational database engine;
- **Vitrage**, the OpenStack Root Cause Analysis service for organizing and handling alarms;
- **Zaqar**, a cloud messaging service with REST API.

Red Hat OpenShift Container Platform is a hybrid cloud PaaS solution built on Kubernetes and RHEL.

OpenShift provides developer services, application services, and platform services e.g. service mesh, pipelines for CI/CD, and full stack logging. It ships packaged with OKD (Origin Kubernetes Distribution).

oc OpenShift CLI. It includes `kubectl`

Kerberos is an authentication protocol that allows hosts to authenticate each other over an insecure network.

The central authentication entity is the **Key Distribution Center (KDC)**, composed of three parts: the database, the Authentication Server, and the Ticket Granting Server.

The **database** stores entries associated with users, hosts, and services. Each entry is called a **principal** and is in the form *username/instance@REALM* (for users) or *service/hostname@REALM* (for services). A **realm** is an authentication administrative domain. A trust relationship between different realms permits cross-authentication i.e. allows users from a realm to authenticate and access the services of another realm.

The **Authentication Server (AS)** replies to the initial authentication request from the client by issuing a **Ticket Granting Ticket (TGT)**.

The **Ticket Granting Server (TGS)** issues service tickets to clients that own a valid TGT. A **ticket** is encrypted with the secret key of the service it is intended for, has a limited validity (10 hours by default), and contains a session key (which is a secret shared between the client and the service). The client will then submit the ticket to an application server in order to prove its identity. Along with the ticket, the client submits an **authenticator** packet containing the user principal and the timestamp, encrypted with the session key.

A client authenticates via Kerberos to an application server through the following steps:

1. The client contacts the AS, making an initial user authentication request
2. The AS replies to the client, sending a TGT (encrypted with the TGS's secret key) and a session key (encrypted with the user's secret key)
3. The client contacts the TGS, sending the TGT and an authenticator (encrypted with the session key)
4. The TGS replies to the client, sending the requested service ticket (encrypted with the service's secret key) and a service session key (encrypted with the session key)
5. The client contacts the application server, sending the service ticket and an authenticator (encrypted with the service session key)

A **keytab** (key table) stores keys for principals. A keytab is usually a file, named as `FILE:/path/krb5.keytab`.

Each entry in a keytab consists of: timestamp, principal name, key version number, encryption type, and encryption key. The keytab file is present in any host that uses Kerberos.

`/etc/krb5/kadm5.keytab`

Keytab file on the KDC

`/etc/krb5/krb5.keytab`

Keytab file on application servers providing kerberized services

User commands:

<code>kinit</code>	Request a TGT for a principal and store it in the credential cache
<code>klist</code>	List principal and tickets contained in the credentials cache, or the keys contained in a keytab file
<code>kswitch</code>	Switch to another credential cache
<code>kdestroy</code>	Destroy the credential cache, deleting all tickets
<code>kvno</code>	Acquire a service ticket for a principal and print out its key version number
<code>kpasswd</code>	Change a principal's password
<code>ksu</code>	Kerberos version of <code>su</code>
<code>krb5-config</code>	Print information useful for compiling and linking programs against the installed Kerberos libraries

Administration commands:

<code>kadmin</code>	Administer a Kerberos system (via <code>kadmin</code>)
<code>kadmin.local</code>	Administer a Kerberos system (via the local KDC database)
<code>kadmind</code>	Start the Kerberos administration server
<code>krb5kdc</code>	Manage the AS and the KDC
<code>kdb5_util</code>	Manage the Kerberos database
<code>kdb5_ldap_util</code>	Manage realms, Kerberos services, and ticket policies
<code>ktutil</code>	Edit a keytab
<code>k5srvutil</code>	Edit keys stored in a keytab
<code>kprop</code>	Propagate the Kerberos database from the primary KDC server to a replica KDC server
<code>kpropd</code>	Listen for and apply updates from <code>kprop</code> . Runs on the replica KDC server
<code>kproplog</code>	Display the log of the Kerberos database updates

Other commands:

<code>sclient</code> <code>sserver</code>	Simple client and server, useful for testing or doing a demo of Kerberos authentication
--	---

RADIUS (Remote Authentication Dial-In User Service) is a networking protocol that provides centralized AAA (Authentication, Authorization, and Accounting) management for users of a network service. It uses TCP/UDP port 1812 for authentication and TCP/UDP port 1813 for accounting.

To access a particular network resource, the user sends a request to a network access server (NAS) which sends a RADIUS **Access Request** message to the RADIUS server. The RADIUS server replies with either an **Access Accept**, an **Access Reject**, or an **Access Challenge** (in case the access requires an additional authentication factor from the client e.g. a PIN, OTP, or card).

Diameter is an improvement over the RADIUS protocol which adds new functionalities and commands.

<code>freeradius</code> <code>radiusd</code>	Run the FreeRADIUS daemon
<code>radclient</code>	Send packets to a RADIUS server
<code>radtest</code>	Send packets to a RADIUS server. Frontend to <code>radclient</code>

The **Name Service Switch (NSS)** is a scheme that allows the local machine to connect and use different name resolution mechanisms e.g. local files, LDAP, DNS, NIS, NIS+. It is used by the OS resolver libraries.

/etc/nsswitch.conf

NSS configuration file.

Each line specifies a database name, followed by the list of possible sources, which will be tried in order to perform name resolution.

```
passwd:    files ldap
shadow:    files
group:     files ldap

hosts:     dns nis nisplus files

ethers:    files nis
netmasks:  files nis
networks:  files nis
protocols: files nis
rpc:       files nis
services:  files nis

automount: files
aliases:   files
```

getent

Get entries from NSS libraries

getent passwd *user*

Get *user*'s password entry

getent group *groupname*

Get entries matching the group *groupname*

getent hosts *host*

Resolve and get the IP address of *host*

Script that performs name resolution on multiple hosts

```
#!/bin/bash

if [ $# -eq 0 ]; then
    echo "Usage: $0 host1 [host2 ... hostN]"
    exit 1
fi

for host in "$@"
do
    ip=$(getent hosts "$host" | awk '{ print $1 }' | head -n1)
    if [ -n "$ip" ]; then
        echo "$host $ip"
    else
        echo "$host [IP not found]"
    fi
done
```

`sssd` is the daemon for the **SSS (System Security Services)** and provides local or remote identity authentication. It is derived from the FreeIPA project.

SSS features its own NSS (Name Service Switch) and PAM (Pluggable Authentication Module) client interfaces, and has its own cache for offline support. Furthermore, it is capable to interface and query different types of directories, databases, and frameworks such as NIS, LDAP, Kerberos, etc.

```
/etc/sss/sssd.conf  
/etc/sss/conf.d/*
```

SSS daemon configuration files

```
sss_cache
```

SSS cache purge utility

```
sss_cache -E
```

Clear the SSS cache and update all records

```
sss_cache -u user
```

Clear the SSS cache for the *user* account only

Identity Management (IdM) is a framework of policies and technologies to ensure that the proper people have access to the proper resources. Similar frameworks are **Identity and Access Management (IAM)** and **IPA (Identity, Policy, and Audit)**.

Single Sign-On (SSO) is an authentication scheme that allows a user to log in to multiple independent services using one set of credentials. It is a subset of **Federated Identity Management (FIdM)**, which handles identity federation i.e. the linking of multiple identities of a user across multiple IdM systems.

SAML (Security Assertion Markup Language) is an XML-based markup language used in authentication procedures, and especially in SSO. SAML exchanges authentication and authorization data between a subject aka principal (i.e. a user), an identity provider, and a service provider. SAML is an open standard build upon XML, HTTP, and SOAP.

The **OAuth** open standard provides secure delegated access i.e. a way for resource owners to authorize third-party applications (consumers) to access their resources from a service provider without disclosing secret credentials; this is done by the means of access tokens. OAuth is designed specifically to operate with HTTP; its latest version is OAuth 2.0.

OIDC (OpenID Connect) is an authentication layer built on top of OAuth 2.0.

Keycloak is an open source IAM and SSO solution, and the upstream project for **Red Hat SSO**.

It supports several standard protocols for authentication and authorization, such as SAML, OAuth 2.0, and OIDC.

Shibboleth is an SSO solution. The Shibboleth Identity Provider and the Shibboleth Service Provider are implementations of SAML.

FreeIPA is an open source IdM system, and the upstream project for **Red Hat Identity Management**.

Its main components are: 389 Directory Server (LDAP server), Dogtag Certificate System (CA), Kerberos, SSSD, NTP, and bind-dyndb-ldap (for integration with DNS).

It features a web interface (Web UI, built as a JavaScript Single Page Application) as well as a CLI (`ipa`).

`ipa`

FreeIPA CLI

The **Dogtag Certificate System** is an open source Certification Authority, written in Java and running on Tomcat.

It is composed of the following six subsystems:

Certificate Authority (CA)	Issues, renews, revokes, and publishes certificates. It also creates and publishes CRLs
Registration Authority (RA)	Authenticates enrollment requests and forwards them to the CA to generate a certificate
Key Recovery Authority (KRA) aka Data Recovery Manager (DRM)	Stores private keys. It can also provide server-side key pair generation
OCSP Manager	Provides OCSP (Online Certificate Status Protocol) functionalities i.e. determine the state of a certificate, and particularly its revocation status, without the need to check a CRL
Token Key Service (TKS)	Manages the master keys used to establish secure channels to the token management system, allowing e.g. smart card tokens to communicate securely with the TPS
Token Processing System (TPS)	Provides RA functionality in the token management system, and establishes secure channels between the client (e.g. smart card management infrastructure) and the backend subsystems (CA, KRA, and TKS)

Git is an open source version control system with a small footprint and very high performances. A Git directory is a complete repository with full history and version tracking abilities, independent of any remote repository. Git commits are identified by a 40-hex-digit hash number, usually shortened to 7 digits, or even less if unambiguous.

<code>git init</code>	Initialize the current directory as a repository
<code>git clone repo</code>	Clone a remote repository. <i>repo</i> can be a URL (SSH, HTTP, HTTPS, FTP, FTPS, Git) or a local path e.g. <code>ssh://user@example.com:8888/path/to/repo.git</code> <code>git://example.com:9999/path/to/repo.git</code> <code>/path/to/repo.git</code>
<code>git checkout branch</code>	Start working into an already existing <i>branch</i>
<code>git checkout -B branch</code>	Create <i>branch</i> and start working into it
<code>git checkout -- file</code>	Discard local changes done to <i>file</i>
<code>git checkout branch file</code>	Copy <i>file</i> from <i>branch</i> to the current branch, and add it to the staging area
<code>git pull</code>	Pull the changes from the remote repository branch to the local branch
<code>git add file</code>	Add <i>file</i> to the staging area (i.e. content staged for the next commit), hence starting to track it
<code>git restore --staged file</code>	Remove <i>file</i> from the staging area, undoing the command <code>git add file</code>
<code>git add .</code>	Add all modified files to the staging area
<code>git rm file</code>	Remove <i>file</i> from the content staged for the next commit
<code>git status</code>	See the status (e.g. files changed but not yet staged) of the current branch
<code>git commit -m "Message"</code>	Commit all staged files in the current branch
<code>git commit -am "Message"</code>	Add all changed files to the staging area in the current branch, and commit them
<code>git merge branch</code>	Merge changes made on <i>branch</i> to the master branch
<code>git push</code>	Push the local commits from the current branch to the remote repository
<code>git push origin branch</code>	Push the local commits from <i>branch</i> to the remote repository
<code>git revert commit</code>	Revert a specific commit
<code>git branch</code>	Show local branches
<code>git branch -r</code>	Show remote branches
<code>git branch -a</code>	Show remote and local branches
<code>git branch -a --contains commit</code>	Show on which branch was done a specific commit number
<code>git branch -d branch</code>	Delete a local branch (which must have been merged in its upstream branch)
<code>git branch -D branch</code>	Delete a local branch (irrespective of its merged status)

<code>git diff</code>	Show the differences between local and remote branch
<code>git diff commit1 commit2</code>	Show the differences between two commits
<code>git diff branch1 branch2</code>	Show the differences between two branches
<code>git diff branch1 branch2 file</code>	Show the differences between two branches for a specific file
<code>git log --all -- file</code>	Show the commits which involved <i>file</i> , across all branches
<code>git log -p --all -S 'string'</code>	Show the commits whose added or deleted lines contain a specific word
<code>git log -p --all -G 'regex'</code>	
<code>git grep string `git show-ref --heads`</code>	Search for <i>string</i> across all branches' heads (i.e. in the latest content only, and not in all the previous commits)
<code>git config --list</code>	Get all currently set options and their values in the Git configuration
<code>git config option</code>	Get the value of <i>option</i>
<code>git config user.name name</code>	Set your username
<code>git config user.email email</code>	Set your email address
<code>onefetch</code> <code>gfetch</code>	Display code information and statistics for a Git repository
<code>glab</code>	GitLab CLI tool
<code>gh</code>	GitHub CLI tool

How to work on a branch different from master (good practice)

1. `git checkout master`
`git pull origin master` Update the local copy of master
2. `git checkout -b branch` Create a new branch from master
3. Edit the files in the local branch, and save the work
4. `git commit -am "Message"` Commit all changes
5. `git push origin branch` Push the branch to the remote repository
6. Create a merge request (on GitHub / GitLab, either via their web GUI or their API CLI) and merge the branch
7. `git checkout master`
`git pull origin master`
`git checkout branch`
`git merge master` While working, keep the local repository in sync with the remote master
8. Repeat from point 3 as necessary
9. `git push origin --delete branch`
`git branch -d branch` Once finished working, delete the remote and local branch to tidy up the repository (optional)

Vagrant is an open source software that allows building and maintaining lightweight and portable virtual environments for software development. It relies on an underlying virtualization or containerization solution e.g. VirtualBox, KVM, VMware, or Docker.

<code>vagrant -h</code>	Print the list of commands recognized by Vagrant
<code>vagrant <i>command</i> -h</code>	Print help about the Vagrant <i>command</i>
<code>vagrant init hashicorp/precise64</code>	Initialize the current directory as a specific Vagrant environment (in this case, Ubuntu 12.04 64-bit) by creating a Vagrantfile on it
<code>vagrant up <i>vmname</i></code>	Start a guest virtual machine and do a first provisioning according to the Vagrantfile
<code>vagrant provision <i>vmname</i></code>	Provision a virtual machine
<code>vagrant ssh <i>vmname</i></code>	Connect via SSH to a virtual machine
<code>vagrant halt <i>vmname</i></code>	Shut down the virtual machine
<code>vagrant destroy <i>vmname</i></code>	Delete the virtual machine and free any resource allocated to it
<code>vagrant status</code>	Print the status of the virtual machines currently managed by Vagrant
<code>vagrant global-status</code>	Print the status of all Vagrant environments on the system, by reading cached data. Completes quickly but results may be outdated
<code>vagrant global-status --prune</code>	Print the status of all Vagrant environments on the system, after rebuilding the environment information cache. Results are always correct but completion takes longer

The directory containing the Vagrantfile on the host machine can be accessed on the guest machine at `/vagrant`.

Ceph is an open source solution for a storage cluster, providing redundancy and high availability, with a single system for access to object storage, block device storage, and file storage.

The components of a **Ceph Storage Cluster** are:

Ceph OSDs (Object Storage Daemons)	Store data and handle data replication, recovery, and rebalancing. At least 3 OSDs are usually required
Ceph Monitors	Maintain maps of the cluster state and handle daemon-client authentication. Uses the Paxos parliament protocol. At least 3 Monitors are usually required
Ceph Managers	Track runtime cluster metrics and allow access to cluster information. At least 2 Managers are usually required
Ceph MDS (Metadata Server)	Stores Ceph File System metadata in order to relieve burden from the cluster. Required when running Ceph File System clients

Ceph is based on RADOS (Reliable Autonomic Distributed Object Store). Storage and retrieval of data are determined by the CRUSH (Controlled Replication Under Scalable Hashing) algorithm, which builds a hierarchical map of the cluster and assigns data to pseudorandomly-chosen OSDs; this improves scalability, reduces performance bottlenecks, and avoids Single Points of Failure.

Ceph File System aka **CephFS** is a POSIX-compliant file system built on top of RADOS.

Ceph Object Storage supports interfaces to the Amazon S3 RESTful API and the OpenStack Swift API. Via the Ceph Object Gateway, it provides a RESTful gateway to a Ceph Storage Cluster.

Ceph Block Device is implemented through RBD (RADOS Block Device) images, which are thin-provisioned and store data striped over multiple OSDs.

<code>ceph-osd</code>	Object Storage daemon for CephFS
<code>ceph-mon</code>	Cluster Monitor daemon for CephFS
<code>ceph-mgr</code>	Manager daemon. It is recommended to place Manager and Monitor daemons in the same nodes
<code>ceph-mds</code>	Metadata Server daemon
<code>ceph</code>	Ceph administration tool for deploying and managing a cluster. CLI tool for Cephadm (since v15.2.0)
<code>ceph-authtool</code>	Management tool for Ceph keyring files (used for authentication)
<code>ceph-volume</code>	Deploy logical volumes as OSDs
<code>ceph-clsinfo</code>	Display information about a specific class object
<code>ceph-run</code>	Restart a daemon
<code>ceph-conf</code>	Display information about Ceph configuration
<code>osdmactool</code>	Manipulate OSD cluster maps, and export or import CRUSH maps
<code>monmaptool</code>	Manipulate Monitor cluster maps
<code>crushtool</code>	Create, compile, decompile, and test CRUSH map files
<code>rados</code>	RADOS utility

Puppet is a software configuration management tool, based on a client-server architecture. It works as follows:

1. A **Puppet agent** (client, running as `root` on each managed node) periodically gathers information (**facts**) about the local node state via the **Facter** system inventory tool.
2. The Puppet agent then communicates this information to the **Puppet master** (server, running as the `puppet` user and listening on TCP port 8140).
3. The Puppet master sends back to the Puppet agent a **catalog** containing the desired configuration for that node.
4. The Puppet agent applies the needed changes (which are idempotent) so that the node's configuration converges with the desired configuration, and sends back a report to the Puppet master.

Puppet configurations are based on **resources** (e.g. "package", "service", "file", "user"). For each resource, a list of **attributes** is specified, with the desired value for each attribute.

Each resource type is implemented through **providers** (e.g. `yum`, `rpm`, `apt`, `opkg` for the resource "package").

Resources managed together as a single unit can be grouped into **classes**; classes are contained in **manifests** which are files with the `.pp` extension.

Modules are directories containing self-contained pieces of configuration and classes for a specific complex setting, e.g. an Apache webserver or a MySQL server.

<code>/etc/puppet/puppet.conf</code>	Configuration file (Open Source Puppet)
<code>/etc/puppetlabs/puppet/puppet.conf</code>	Configuration file (Puppet Enterprise)
<code>facter</code>	Gather the facts about the managed node, returning a list of key-value pairs
<code>puppet agent</code>	Main Puppet client. Retrieves the node's desired configuration from the Puppet master and applies it
<code>puppet agent --enable</code>	Enable the Puppet agent on the node
<code>puppet agent --disable "Reason for disabling"</code>	Disable the Puppet agent on the node
<code>cat \$(puppet config print vardir)/state/agent_disabled.lock</code>	Print the reason why the Puppet agent is currently disabled. If the Puppet agent is enabled instead, this lockfile does not exist
<code>puppet agent --noop</code>	Perform a dry run, displaying the changes that Puppet would have applied without actually applying them
<code>puppet --version</code> <code>puppet agent --version</code> <code>puppet master --version</code>	Show version of different Puppet components
<code>puppet module list</code>	List all modules installed in Puppet
<code>puppet resource user username</code>	Inspect the state of the resource "user" with respect to <i>username</i>
<code>puppet resource service httpd enable=false</code>	Modify the state of the resource "service" (in this case, disable the HTTP server)
<code>puppet describe user</code>	Show information about the resource "user"
<code>puppet describe --list</code>	List all resource types
<code>puppet describe user --providers</code>	Return the list of providers for the resource "user"
<code>puppet apply modulename/init.pp</code>	Apply a manifest one time only
<code>puppet cert operation</code>	Manage the SSL certificates used for communications between master and agents
<code>puppet master --configprint basemodulepath</code>	Display the specified configuration value

Manifest and other related files

basemodulepath/modulepath/ssh/manifests/init.pp

Default class for the SSH service. Ensures that SSH is installed and running

```
class ssh {
  package { ['ssh']:
    ensure => present,
    name => $::operatingsystem ? {
      'Redhat' => 'openssh',
      'Ubuntu' => 'ssh',
      default  => 'openssh',
    },
  }
  file { ['/etc/ssh/sshd_config']:
    ensure => file,
    owner  => 'root',
    group  => 'root',
    mode   => '0644',
    require => Package['openssh'],
    source => 'puppet:///modules/ssh/sshd_config',
  }
  service { ['sshd']:
    ensure => running,
    enable => true,
    require => File['/etc/ssh/sshd_config'],
  }
  notify { ['Test message for SSH class': ]
}
```

basemodulepath/environment/hiera.yaml

Hiera configuration file

```
---
ntp::servers:
  - time.example.org
  - 0.pool.ntp.org
```

basemodulepath/modulepath/templates/host.erb

ERB (Embedded Ruby) template. Allows embedding dynamic variables that can be resolved in the calling manifest

The IP address is <%= @ipaddress %>

Node definition. Assigns one or more classes to a node

```
node 'n1.example.org' {
  include ssh
  include apache
}
```

Resource relationship	
<pre>package { 'openssh-server': ensure => present, before => File['/etc/ssh/sshd_config'], } or file { '/etc/ssh/sshd_config': ensure => file, mode => '0600', source => 'puppet:///modules/sshd/sshd_config', require => Package['openssh-server'], }</pre>	Ensures that the SSH server package is installed before the SSH daemon configuration file
<pre>file { '/etc/ssh/sshd_config': ensure => file, mode => '0600', source => 'puppet:///modules/sshd/sshd_config', notify => Service['sshd'], } or service { 'sshd': ensure => running, enable => true, subscribe => File['/etc/ssh/sshd_config'], }</pre>	Notifies the SSH service (restarting it) any time the SSH daemon configuration file is changed

Ansible is an open source tool (made by Red Hat) for configuration management, software provisioning, and application deployment. It is agentless and connects to the managed machines via SSH pubkey authentication, requiring only OpenSSH and Python to be installed on the managed nodes.

The configuration for a managed node is specified in a **playbook**, written in YAML. A playbook contains one or more **plays** to be run in sequence, each of which lists a number of small specific **tasks** to be run in sequence. Each task calls a **module**, which describes the desired state of the system and executes the operation to do so (e.g. start the Apache server, copy a file, verify that a package is installed, rename a database table). A module should be idempotent i.e. it should set the system always in the same state regardless of how many times it is run.

Handlers are tasks that are run only when notified (via the command `notify`), when a change is made on a machine. The **inventory** lists and categorizes all hosts on which tasks have to be executed. It is also possible to define **roles** to categorize hosts and automatically load variables, tasks, handlers, and other artifacts pertaining to that role. When a playbook is run, first it collects system and environment information (**facts**) which is then stored in multiple variables named `ansible_varname`.

Ansible uses the Jinja template engine to enable access to tests, filters, variables, and dynamic execution. Also available is a set of plugins that expand Ansible's core functionalities.

<code>/etc/ansible/hosts</code>	Inventory file, containing the list of hosts managed by Ansible. Can be in INI or YAML format
<code>ansible hosts -m module -a options</code>	Run the <i>module</i> with the specific <i>options</i> on the <i>hosts</i>
<code>ansible hosts -m ansible.builtin.copy \</code> <code>-a "src=/path/to/file dest=/tmp/"</code>	Copy a file to the <i>hosts</i>
<code>ansible hosts -m ansible.builtin.yum \</code> <code>-a "name=httpd state=latest"</code>	Ensure that <i>hosts</i> have the <code>httpd</code> package installed and updated to the latest version
<code>ansible hosts -a "/sbin/shutdown"</code>	Shutdown the <i>hosts</i>
<code>ansible all -m ping</code>	Ping all hosts in the inventory (by executing the "ping" module)
<code>ansible all -m ansible.builtin.setup</code>	Show all facts
<code>ansible-playbook playbook.yml</code>	Apply the specified playbook
<code>ansible-lint playbook.yml</code>	Check the syntax of the specified playbook
<code>ansible-pull</code>	Pull a remote copy of Ansible on each managed node and download the playbooks from a source repository. This inverts the default push architecture of Ansible into a pull architecture

ansible command line options	
<code>-m module</code>	Run the specified <i>module</i> . If not specified, Ansible uses the default module "ansible.builtin.command"
<code>-a options</code>	Apply the specified <i>options</i> when running the module
<code>-f n</code>	Fork <i>n</i> processes when running the module. If not specified, default is 5
<code>-u user</code>	Run the module as <i>user</i> . If not specified, default is current user
<code>--become</code>	Run the module as root
<code>--become --ask-become-pass</code>	Run the module as root, asking for the password

Inventory and variables files

hosts Inventory file, defining two groups <pre>[webservers] 10.0.1.17 10.0.1.18 [dbservers] 10.0.2.42</pre>	group_vars/all Variables applied to all host groups <pre>--- httpd_port: 80 ntpserver: 192.168.0.13 repository: https://foobar.org/repo/website.git</pre>	group_vars/dbservers Variables applied to hosts in the "dbservers" group <pre>--- mysqldservice: mysqld mysql_port: 3306 dbuser: jdoe dbname: mydb dbpassword: mys3cr3t</pre>
--	---	---

Main playbook file

```
lamp.yml
Playbook that deploys the whole LAMP stack on the site

---
- name: Apply common configuration to all nodes
  hosts: all
  remote_user: root

  roles:
    - common

- name: Deploy and configure webservers
  hosts: webservers
  remote_user: root

  roles:
    - web

- name: Deploy MySQL and configure databases
  hosts: dbservers
  remote_user: root

  roles:
    - db
```

Files for role "common"

roles/common/handlers/main.yml Handler for general notifications, called from other plays <pre>--- - name: Restart NTP service: name: ntpd state: restarted</pre>	roles/common/tasks/main.yml General play run on all nodes <pre>--- - name: Install NTP yum: name: ntp state: present tags: ntp - name: Configure NTP template: src: ntp.conf.j2 dest: /etc/ntp.conf tags: ntp notify: Restart NTP - name: Start NTP service: name: ntpd state: started enabled: yes tags: ntp</pre>	roles/common/templates/ntp.conf.j2 Jinja template for NTP configuration file <pre>driftfile /var/lib/ntp/drift restrict 127.0.0.1 server {{ ntpserver }} includefile /etc/ntp/crypto/pw keys /etc/ntp/keys</pre>
---	---	--

Files for role "db"	
<pre>roles/db/handlers/main.yml Handler for DB tier notifications --- - name: Restart MySQL service: name: mysqld state: restarted - name: Restart iptables service: name: iptables state: restarted</pre>	<pre>roles/db/tasks/main.yml Install MySQL, then create database and database user --- - name: Install the MySQL package yum: name: "{{ item }}" state: installed with_items: - mysql-server - MySQL-python - name: Configure MySQL template: src: my.cnf.j2 dest: /etc/my.cnf notify: Restart MySQL - name: Start MySQL service: name: mysqld state: started enabled: yes - name: Insert iptables rule for MySQL lineinfile: dest: /etc/sysconfig/iptables state: present regexp: "{{ mysql_port }}" insertafter: "^:OUTPUT " line: "-A INPUT -p tcp --dport {{ mysql_port }} -j ACCEPT" notify: Restart iptables - name: Create database mysql_db: name: "{{ dbname }}" state: present - name: Create database user and set password mysql_user: name: "{{ dbuser }}" password: "{{ dbpassword }}" priv: " *.*:ALL" host: '%' state: present</pre>
<pre>roles/db/templates/my.cnf.j2 Jinja template for MySQL configuration file [mysqld] datadir=/var/lib/mysql socket=/var/lib/mysql/mysql.sock user=mysql symbolic-links=0 port={{ mysql_port }} [mysqld_safe] log-error=/var/log/mysqld.log pid-file=/var/run/mysqld/mysqld.pid</pre>	

Files for role "web"		
roles/web/handlers/main.yml Handler for web tier notifications	roles/web/tasks/main.yml	Main task, calls the other two playbooks
<pre> --- - name: Restart iptables service: name: iptables state: restarted </pre>	<pre> --- - include: install_httpd.yml - include: copy_website.yml </pre>	
	roles/web/tasks/copy_website.yml	Copy the code from the git repository
	<pre> --- - name: Copy website from repo git: repo: "{{ repository }}" dest: /var/www/html/ - name: Create the index.php file template: src: index.php.j2 dest: /var/www/html/index.php </pre>	
	roles/web/tasks/install_httpd.yml	Install HTTP, PHP, and Git modules
	<pre> --- - name: Install httpd packages yum: name: "{{ item }}" state: present with_items: - httpd - php - php-mysql - git - name: Insert iptables rule for httpd lineinfile: dest: /etc/sysconfig/iptables create: yes state: present regexp: "{{ httpd_port }}" insertafter: "^:OUTPUT " line: "-A INPUT -p tcp --dport {{ httpd_port }} -j ACCEPT" notify: Restart iptables - name: Check that httpd is running service: name: httpd state: started enabled: yes </pre>	
roles/web/templates/index.php.j2 Jinja template for the website root file index.php	<pre> <html> <head> <title>LAMP stack and website deployed via Ansible</title> </head> <body> Homepage
 <?php echo "Hostname: " . exec('hostname') . "
"; echo "Database list:
"; {% for host in groups['dbservers'] %} \$link = mysqli_connect('{{ hostvars[host].ansible_default_ipv4.address }}', '{{ hostvars[host].dbuser }}', '{{ hostvars[host].dbpassword }}') or die(mysqli_connect_error(\$link)); {% endfor %} while (\$r = mysqli_fetch_assoc(mysqli_query(\$link, "SHOW DATABASES;")) {echo \$r['Database'] . "\n";} ?> </body> </html> </pre>	

Ansible Tower is a framework that allows implementing enterprise-wide automation web console for centralized configuration and management of Ansible elements. It provides web services, a REST API, RBAC, job scheduling, Tower clustering, and graphical inventory management. Its upstream open source project is **AWX**.

Ansible Tower is part of the **Red Hat Ansible Automation Platform**.

<code>tower-cli</code>	CLI for Ansible Tower (legacy)
<code>awx-cli</code>	

<code>awx</code>	CLI for Ansible Tower
------------------	-----------------------

Tag		Attributes	
<h1> ... <h6> Heading		<code>align=left center right justify</code>	Heading alignment †

 Line break	Line break and carriage return		
<hr> Horizontal line		<code>align=left center right</code> <code>noshade</code> <code>size=npixels</code> <code>width=npixels percent%</code>	Line alignment † Solid rendering instead of 3D † Line height Line width
<p> Paragraph <div> Section		<code>align=left center right justify</code>	Paragraph or section alignment †
 Group	Group of elements		
<a> Anchor	Hyperlink	<code>charset=encoding</code> <code>coords=left,top,right,bottom cx,cy,radius x1,y1,...,xn,yn</code> <code>href=url</code> <code>hreflang=language</code> <code>name=section</code> <code>rel rev=alternate stylesheet start next prev contents index glossary copyright chapter section subsection appendix help bookmark</code> <code>shape=rectangle circle polygon</code> <code>target=_blank _parent _self _top</code> <code>type=mimetype</code>	Character encoding of target URL Coordinates of region; depends on shape Target URL for the link Language of document at the target URL Name of anchor for document bookmarking Relationship between this document and the target URL (<code>rel</code>) or vice versa (<code>rev</code>) Shape of region Destination of target URL MIME type of target URL
<dl> Definition list			
<dt> Definition term			
<dd> Definition description	Description of a definition term		
 Ordered list		<code>compact=compact</code> <code>start=firstnumber</code> <code>type=A a I i 1</code>	List must be more compact † Number to start the list on † List numbers type †
 Unordered list		<code>compact=compact</code> <code>type=disc square circle</code>	List must be more compact † List type †
 List item		<code>type=disc square circle A a I i 1</code> <code>value=itemno</code>	List item type † List item value †

† = deprecated

Tag	Attributes		
<i> Italic			
 Bold			
<s> Strike-through	Strike-through text †		
<u> Underlined	Underlined text †		
<big> Bigger			
<small> Smaller			
<sub> Subscript			
<sup> Superscript			
<tt> Teletype	Monospaced text		
 Emphasized			
 Strong			
 Deleted	Deleted/inserted text	<code>cite=url</code>	URL to document explaining deletion/insertion
<ins> Inserted		<code>datetime=yyyy-mm-dd</code>	When the text was deleted/inserted
<pre> Preformatted		<code>width=ncharacters</code>	Max number of characters per line †
<code> Code	Source code text		
<samp> Sample	Sample code text		
<kbd> Keyboard	Keyboard key		
<var> Variable	Variable name		
<cite> Citation	Citation block		
<blockquote> Quotation		<code>cite=url</code>	URL to document containing the quote
<q> Short quotation			
<address> Address	Address block		
<abbr> Abbreviation			
<acronym> Acronym			
<dfn> Definition	Definition term		
 Font	Font †	<code>color=rgb(r,g,b) #rrggbb color</code>	Text color
		<code>face=fontname</code>	Text font
		<code>size=[1 ... 7] [-6 ... +6]</code>	Text size
<bdo> Bidirectional override		<code>dir=ltr rtl</code>	Direction of text: left-to-right or right-to-left
<xmp> XMP	Non-formatted text † (ignores other HTML tags)		
other tags	Attributes common to almost all other tags	<code>class=class style</code>	Class of the element
		<code>id=id</code>	Unique ID of the element
		<code>style=styledef</code>	Inline style definition
		<code>title=tooltip</code>	Text of the tooltip to display
		<code>dir=ltr rtl</code>	Direction of text: left-to-right or right-to-left
		<code>lang=language</code>	Language of the content
		<code>accesskey=character</code>	Keyboard shortcut for the element
		<code>tabindex=ntab</code>	N of tab for the element

† = deprecated

Tag	Attributes	
 Image	<code>align=top bottom left middle right</code>	Image alignment with respect to surrounding text †
	<code>alt=alternatetext</code>	Description of the image for text-only browsers
	<code>border=npixels</code>	Border width around the image †
	<code>height=npixels percent%</code>	Image height
	<code>hspace=npixels</code>	Blank space on the left and right side of image †
	<code>ismap=url</code>	URL for server-side image map
	<code>longdesc=url</code>	URL containing a long description of the image
	<code>src=url</code>	URL of the image
	<code>usemap=url</code>	URL for client-side image map
	<code>vspace=npixels</code>	Blank space on top and bottom of image †
	<code>width=npixels percent%</code>	Image width
<map> Image map	<code>id=id</code>	Unique ID for the map tag
	<code>name=name</code>	Unique name for the map tag
<area> Area of image map	<code>alt=alternatetext</code>	Description of area for text-only browsers
	<code>coords=left,top,right,bottom cx,cy,radius x1,y1,...,xn,yn</code>	Coordinates of clickable area; depends on <code>shape</code>
	<code>href=url</code>	Target URL of area
	<code>nohref=true false</code>	Excludes or includes the area from image map
	<code>shape=rectangle circle polygon</code>	Shape of area
	<code>target=_blank _parent _self _top</code>	Destination of target URL

† = deprecated

Tag	Attributes	
<table> Table	<code>align=left center right</code> <code>bgcolor=rgb(<i>r,g,b</i>) #<i>rrggbb</i> color</code> <code>border=<i>npixels</i></code> <code>cellpadding=<i>npixels percent%</i></code> <code>cellspacing=<i>npixels percent%</i></code> <code>frame=void above below lhs rhs hsides vsides box border</code> <code>rules=none groups rows cols all</code> <code>summary=<i>summary</i></code> <code>width=<i>npixels percent%</i></code>	Table alignment † Table background color † Border width Space around the content of each cell Space between cells Visibility of sides of the table border Horizontal or vertical divider lines Summary of the table for text-only browsers Table width
<tr> Table row	<code>align=left center right justify char</code> <code>bgcolor=rgb(<i>r,g,b</i>) #<i>rrggbb</i> color</code> <code>char=<i>character</i></code> <code>charoff=<i>npixels percent%</i></code> <code>valign=top middle bottom baseline</code>	Horizontal text alignment Row background color † Character to align text on, if <code>align=char</code> Alignment offset to first character, if <code>align=char</code> Vertical text alignment
<td> Table cell <th> Table header	<code>abbr=<i>content</i></code> <code>align=left center right justify char</code> <code>axis=<i>category</i></code> <code>bgcolor=rgb(<i>r,g,b</i>) #<i>rrggbb</i> color</code> <code>char=<i>character</i></code> <code>charoff=<i>npixels percent%</i></code> <code>colspan=<i>ncolumns</i></code> <code>headers=<i>headerid</i></code> <code>height=<i>npixels</i></code> <code>nowrap</code> <code>rowspan=<i>nrows</i></code> <code>scope=col colgroup row rowgroup</code> <code>valign=top middle bottom baseline</code> <code>width=<i>npixels percent%</i></code>	Abbreviated content in a cell Horizontal text alignment Cell name Cell background color † Character to align text on, if <code>align=char</code> Alignment offset to first character, if <code>align=char</code> Number of columns this cell spans on Cell header information for text-only browsers Cell height † Text in cell stays on a single line † Number of rows this cell spans on Target for cell header information Vertical text alignment Cell width †

† = deprecated

Dec	Hex	Char	Dec	Hex	Char	Dec	Hex	Char	Dec	Hex	Char
0	0	NUL	32	20	space	64	40	@	96	60	`
1	1	SOH	33	21	!	65	41	A	97	61	a
2	2	STX	34	22	"	66	42	B	98	62	b
3	3	ETX	35	23	#	67	43	C	99	63	c
4	4	EOT	36	24	\$	68	44	D	100	64	d
5	5	ENQ	37	25	%	69	45	E	101	65	e
6	6	ACK	38	26	&	70	46	F	102	66	f
7	7	BEL	39	27	'	71	47	G	103	67	g
8	8	BS	40	28	(	72	48	H	104	68	h
9	9	TAB	41	29	)	73	49	I	105	69	i
10	A	LF	42	2A	*	74	4A	J	106	6A	j
11	B	VT	43	2B	+	75	4B	K	107	6B	k
12	C	FF	44	2C	,	76	4C	L	108	6C	l
13	D	CR	45	2D	-	77	4D	M	109	6D	m
14	E	SO	46	2E	.	78	4E	N	110	6E	n
15	F	SI	47	2F	/	79	4F	O	111	6F	o
16	10	DLE	48	30	0	80	50	P	112	70	p
17	11	DC1	49	31	1	81	51	Q	113	71	q
18	12	DC2	50	32	2	82	52	R	114	72	r
19	13	DC3	51	33	3	83	53	S	115	73	s
20	14	DC4	52	34	4	84	54	T	116	74	t
21	15	NAK	53	35	5	85	55	U	117	75	u
22	16	SYN	54	36	6	86	56	V	118	76	v
23	17	ETB	55	37	7	87	57	W	119	77	w
24	18	CAN	56	38	8	88	58	X	120	78	x
25	19	EM	57	39	9	89	59	Y	121	79	y
26	1A	SUB	58	3A	:	90	5A	Z	122	7A	z
27	1B	ESC	59	3B	;	91	5B	[	123	7B	{
28	1C	FS	60	3C	<	92	5C	\	124	7C	
29	1D	GS	61	3D	=	93	5D	]	125	7D	}
30	1E	RS	62	3E	>	94	5E	^	126	7E	~
31	1F	US	63	3F	?	95	5F	_	127	7F	DEL Delete

Characters 0-31 and 127 are non-printable.

`ascii` Display an ASCII table
`man ascii`

`showkey -a` Prompt for pressing a key and display its ASCII value in decimal, octal, and hex

Confidentiality, Integrity, and Availability (aka the **CIA triad**) are the basic policies of Information Security. Confidentiality ensures that access to information is limited to people and groups with the correct rights, integrity ensures that information has not been improperly modified, and availability ensures that a system is operable and functioning.

Access control types	
Discretionary Access Control (DAC) aka need-to-know	Allows the user that has access to the resource to decide with whom to share it. File access is regulated by user and group permissions. In Linux, this is the standard access model.
Mandatory Access Control (MAC)	A particular user can access a resource only if they have been given explicit access right to it. The end user is not allowed to choose who can access the resource, or to pass privileges. In Linux, this is implemented via SELinux.
Role-Based Access Control (RBAC)	Access permissions are based on the access policies determined by the system. Users are assigned access to resources on a one-to-one basis.
Rule-Based Access Control (RuBAC)	Endpoint devices (e.g. firewalls) verify the requests to access network resources against a set of rules based on IP addresses, port numbers, etc.

A **threat** is an entity, circumstance, or event with the potential to adversely impact a computer system through unauthorized access, destruction, disclosure, modification of data, or Denial of Service.

A **vulnerability** is a weakness in a computer system that could be exploited by a threat source.

A **risk** is the probability that a particular security threat will exploit a system vulnerability, according to the risk equation: $\text{risk} = \text{threat} \times \text{vulnerability}$.

An **attack** is an attempt to gain unauthorized access to a computer system's services, resources, or information; can also be considered an attempt to compromise a computer system's confidentiality, integrity, or availability.

Vulnerability management life cycle	
Creation of a baseline ↓	Identify and prioritize critical assets
Vulnerability assessment ↓	Identify and prioritize vulnerabilities. Examine the abilities of a system, applications, security procedures, and controls, to withstand an attack
Risk assessment ↓	Analyze and evaluate risks in order to determine possible incidents, their likelihood, their consequences, and the tolerance of each critical asset for such events. The components of risk assessment are technical safeguards , organizational safeguards , physical safeguards , and administrative safeguards
Remediation ↓	Reduce the severity of vulnerabilities (via action plans, patches, hotfixes, etc.)
Verification ↓	Verify all the previous phases (via scanners, reports, etc.)
Monitor	Monitor regularly the system to maintain the required level of security

The **Common Vulnerability Scoring System (CVSS)** assessment consists of three metrics for measuring vulnerabilities: **base metrics**, **temporal metrics**, and **environmental metrics**. Each metric sets a severity score from 1 to 10.

Common Vulnerabilities and Exposures (CVE) is a public list of identifiers for software vulnerabilities.

The **Metasploit Framework** is a well-known tool and the de facto standard for developing and exploiting security vulnerabilities in systems and applications. The GUI equivalent is **Armitage**.

Metasploit has a modular architecture. **Exploit** modules are the basic modules used to encapsulate an exploit.

Payload modules can be **singles** (self-contained), **stagers** (modules that establish a network connection to the attacked system), and **stages** (downloaded by stagers).

Auxiliary modules are used to perform one-time actions such as port scanning, DoS, or fuzzing.

NOP modules generate no-op instructions (0x90 for x86 microchips) used to keep the payload size consistent by padding out buffers.

The steps for carrying out an attack are, in order: configuring an exploit, setting up the exploit options, selecting a target, selecting a payload, and launching the exploit.

msfconsole	Launch the Metasploit Framework CLI
msfupdate	Update the Metasploit Framework
msfvenom	Generate and encode a payload for an exploit. Replaces the old <code>msfpayload</code> and <code>msfencode</code> tools

Metasploit base commands	
help	Show the list of Metasploit commands
help <i>command</i>	Show detailed help about <i>command</i>
db_status	Check database connection status
creds	Display all credentials in the database
use <i>module</i>	Load and use <i>module</i>
setg <i>variable value</i>	Assign <i>value</i> to a global <i>variable</i>
getg <i>variable</i>	Get the value of a global <i>variable</i>
unsetg <i>variable</i>	Unset the value of a global <i>variable</i>
connect <i>host port</i>	Connect to a remote <i>host</i> on <i>port</i>
sessions	Display information about active sessions
threads	Display information about background threads and manipulate them
banner	Display a Metasploit banner
history	Show Metasploit command history

Metasploit module commands	
info	See information about the currently loaded module
show payloads	Show the list of compatible payloads for the currently loaded module
show options	Show all options (variables) available for the module, along with their descriptions and set values
set <i>variable value</i>	Assign <i>value</i> to a context-specific <i>variable</i>
get <i>variable</i>	Get the value of a context-specific <i>variable</i>
unset <i>variable</i>	Unset the value of a context-specific <i>variable</i>
check <i>host</i>	Check if <i>host</i> is vulnerable
reload	Reload the module
rexploit rerun	Reload and execute the module
exploit run	Execute the module

Meterpreter is an advanced and dynamically extensible payload for a Metasploit attack that provides the attacker with an interactive shell (Meterpreter session) to the target remote machine. This session is obtained by running from Metasploit an exploit with the appropriate payload e.g. `php/meterpreter/reverse_tcp`. The payload is deployed using in-memory DLL injection.

Meterpreter commands	
<code>help</code>	Show a list of all Meterpreter commands
<code>background</code>	Send the Meterpreter session to background and return to the MSF CLI
<code>cat file</code>	Display the contents of <i>file</i>
<code>edit file</code>	Edit a text <i>file</i> (via Vim)
<code>ls</code>	List files on the target machine
<code>ps</code>	Display processes on the target machine
<code>download file</code>	Download <i>file</i> from the target machine
<code>upload file path</code>	Upload <i>file</i> from the local machine to the <i>path</i> on the target machine
<code>execute -f command</code>	Execute a <i>command</i> on the target machine
<code>resource file</code>	Execute on the target machine the Meterpreter commands listed in the local <i>file</i>
<code>shell</code>	Run an OS shell on the target machine
<code>sysinfo</code>	Get information about the target machine: OS, machine name, etc.
<code>ipconfig</code>	Display network configuration of the target machine
<code>localtime</code>	Display date and time of the target machine
<code>clearev</code>	On a MS Windows target machine, clear all Application, Security, and System logs
<code>webcam_list</code>	List all operative webcams on the target machine
<code>webcam_snap</code>	Take a snapshot from the webcam on the target machine
<code>quit</code> <code>exit</code>	End the Meterpreter session

Aircrack-ng is a suite of tools for Wi-Fi security. It includes utilities for wireless network sniffing, attack, key cracking, and testing.

`aircrack-ng options file`

Key cracker. Cracks WEP or WPA/WPA2-PSK keys from a capture *file* (in `.cap` or `.ivs` format). Possible *options* are:

- `-a n` Attack mode ($n=1$ for WEP, $n=2$ for WPA/WPA2-PSK)
- `-e essid` Specify the Access Point to use
- `-K` Use the Korek WEP cracking technique
- `-z` Use the PTW WEP cracking technique
- `-k n` Disable the Korek WEP attack number n (where n is between 1 and 17)
- `-n len` Specify WEP key length
- `-s` Show WEP key in ASCII while cracking
- `-w file` Wordlist file to use for WEP or WPA/WPA2 key cracking

`aireplay-ng attack options`

Packet injector. Replays packets to perform an attack, where *attack* is one of:

- `-0` Deauthentication attack
- `-1` Fake authentication attack
- `-2` Interactive packet replay attack
- `-3` ARP Request replay attack
- `-4` Chopchop attack
- `-5` Fragmentation attack (against WEP)
- `-6` Caffè Latte attack
- `-7` Hirte attack
- `-8` WPA Migration Mode attack
- `-9` Injection test

`airodump-ng options interface`

Packet sniffer. Captures packets by listening to the network *interface*. Possible *options* are:

- `--ivs` Save only captured IVs
- `-w file` Write sniffed packets in a capture *file*
- `-o format` Use *format* for the capture file: can be `pcap`, `ivs`, `csv`, `gps`, `kismet`, `netxml`, or `logcsv`

<code>airoscript-ng</code>	User-friendly interface for <code>aircrack-ng</code>
<code>airgraph-ng</code>	Tool to generate graphs of relationships between wireless devices, using data from <code>airodump-ng</code>
<code>packetforge-ng</code>	Generator of encrypted packets for injection
<code>airbase-ng</code>	Tool for attacks against wireless clients (and not Access Points)
<code>airserv-ng</code>	Tool to access the wireless NIC from other machines
<code>tkiptun-ng</code>	WPA-TKIP attack tool
<code>wesside-ng</code>	WEP key recovery tool
<code>airdecloak-ng</code>	Tool to remove WEP cloaking from PCAP dump files
<code>airolib-ng</code>	Manager for ESSID and password lists. These are used for WPA and WPA2 cracking
<code>airdecap-ng</code>	Decryption tool for WEP and WPA
<code>airmon-ng</code>	Tool to set up wireless NICs in monitor mode
<code>airtun-ng</code>	Virtual tunnel interface creator
<code>airdriver-ng</code>	Tool that displays information about wireless drivers on the system
<code>airdrop-ng</code>	Tool to force deauthentication of users
<code>ivstools</code>	Tool to extract or merge IVs from a PCAP dump file
<code>kstats</code>	Program that displays statistical FMS algorithm votes for an IVs dump with a specific WEP key
<code>makeivs</code>	Tool to generate a dummy IVs dump file with a specific WEP key
<code>versuck-ng</code>	Tool to calculate the default WEP key for Verizon Actiontec wireless routers
<code>easside-ng</code>	Tool to communicate with an Access Point without knowing its WEP key
<code>buddy-ng</code>	Helper server for <code>easside-ng</code>

How to crack WEP

1. `airmon-ng start wlan0` Start the wireless NIC into monitor mode on the same channel as the AP, and test injection capabilities to the AP
2. `airodump-ng --ivs -w capture wlan0` Discover the list of active wireless machines. Note ESSID and BSSID of the target AP (let us assume they are respectively *ap_essid* and *ap_mac*). Keep this command running to capture the generated IVs
3. `aireplay-ng -1 0 -e ap_essid \ -a ap_mac -h mac wlan0` Do a fake authentication with the target AP using your NIC (with MAC address *mac*)
4. `aireplay-ng -3 -b ap_mac -h mac wlan0` To capture a large number of IVs in a short time, run `aireplay-ng` in ARP Request replay mode
5. `aircrack-ng -s capture.ivs` Once `airodump-ng` has captured at least 50000 IVs, crack the WEP key

How to crack WPA-PSK or WPA2-PSK

1. `airmon-ng start wlan0` Put the wireless NIC into monitor mode
2. `airodump-ng -w capture wlan0` Discover the list of active wireless machines. Note BSSID of the target client (let us assume it's *cl_mac*). Keep this command running
3. `aireplay-ng --deauth 11 -a cl_mac` Deauthenticate the client. The client will try to authenticate again, and `airodump-ng` will capture the authentication packet sent during the WPA/WPA2 four-way handshake
4. `aircrack-ng -a 2 -w capture.cap` Analyze the capture dump to crack the WPA/WPA2 key

The only way to crack WPA/WPA2 is to sniff the Pairwise Master Key associated with the four-way handshake authentication process. Therefore it is important to choose a complex WPA/WPA2 random password at least 20 characters long.

How to decloak a hidden SSID

1. `airmon-ng start wlan0` Put the wireless NIC into monitor mode
2. `airodump-ng --ivs -w capture wlan0` Discover the list of active wireless clients. Find the one where the ESSID is hidden (it shows only its string length) and note its BSSID (let us assume it's *cl_mac*). Keep this command running
3. `aireplay-ng --deauth 11 -a cl_mac` Deauthenticate the client
4. The output of `airodump-ng` will now show the hidden SSID

How to perform a MitM attack

1. `airmon-ng start wlan0` Put the wireless NIC into monitor mode
2. `airodump-ng --ivs -w capture wlan0` Discover the list of active wireless clients. Note BSSID of the target client (let us assume it's *cl_mac*). Note ESSID and BSSID of the AP you are currently accessing (let us assume they are respectively *ap_essid* and *ap_mac*). Keep this command running
3. `aireplay-ng --deauth 11 -a cl_mac` Deauthenticate the client
4. `aireplay-ng -1 0 -e ap_essid \ -a ap_mac -h cl_mac wlan0` Perform a fake authentication attack, associating your NIC with the AP you are currently accessing

The **Firmware Mod Kit** is a toolkit to extract, deconstruct, modify, rebuild, and flash firmware images for Linux-based routers, IoT devices, embedded devices, and most other devices that use common firmware formats and filesystems such as TRX/uImage and SquashFS/CramFS.

It can be used by an attacker after compromising a device, to maintain access.

<code>extract-firmware.sh</code>	Extract the firmware
<code>build-firmware.sh</code>	Rebuild the firmware
<code>ddwrt-gui-extract.sh</code>	Extract Web GUI files from extracted DD-WRT firmware
<code>ddwrt-gui-rebuild.sh</code>	Restore modified Web GUI files to extracted DD-WRT firmware

ClamAV is an open source antivirus solution, created by Cisco. It is designed to scan files quickly.

<code>clamd</code>	ClamAV scanning daemon
<code>clamonacc</code>	ClamAV on-access daemon, providing real-time protection. Registers for file access notifications from the Linux kernel, and in response it submits scans to <code>clamd</code>
<code>clamconf</code>	ClamAV configuration tool
<code>clamconf -g freshclam.conf</code>	Generate an example configuration file
<code>freshclam</code>	ClamAV virus database updater
<code>clamscan</code>	ClamAV tool to scan files for viruses
<code>clamscan -v file</code>	Scan a file for viruses
<code>clamscan -v -r dir</code>	Scan a directory for viruses
<code>setsebool -P antivirus_can_scan_system 1</code>	Authorize antivirus scans on systems with SELinux enabled

/etc/crontab for ClamAV

```
03 0 * * * root /usr/bin/freshclam --quiet
27 0 * * * root /usr/bin/clamscan -i --quiet -l /var/log/clamav/scan-`date "+%F"`.log -r /
```

This crontab updates the virus database every day at 00:03, and performs a full system scan every day at 00:27.

Microsoft Defender for Endpoint (MDE) on Linux is an enterprise security platform to protect against information security threats. As a cloud-native solution, it uses Azure Arc as the management control plane.

<code>systemctl status mdatp</code>	Check Defender service status
<code>mdatp health</code>	Display information about the Defender status
<code>mdatp health --field engine_load_status</code>	Display information about a field of the Defender status
<code>mdatp connectivity test</code>	Check network connectivity
<code>mdatp definitions update</code>	Update Defender AV database
<code>mdatp config proxy set --value http://proxy:6666</code>	Configure Defender to use <i>proxy</i> as gateway
<code>mdatp config proxy reset</code>	Erase Defender's proxy configuration
<code>mdatp edr tag set --name GROUP --value MDE-Management</code>	Set the tag MDE-Management
<code>azcmagent</code>	Azure Connected Machine agent CLI tool

<code>arpspoof</code>	ARP spoofing tool
<code>arpoison</code>	ARP cache update utility. Can be used to craft custom ARP packets
<code>arpstraw</code>	ARP spoofing detection tool
<code>arpon</code>	ARP handler inspector. Useful to detect ARP spoofing
<code>arpwatch</code>	Tracker of MAC to IP address pairings. Useful to detect ARP spoofing
<code>ettercap</code>	Network security tool for ARP poisoning and man-in-the-middle attacks over the LAN
<code>macchanger</code>	Tool to perform MAC spoofing. Changes the NIC to a random MAC address
<code>macof</code>	DoS tool for MAC flooding
<code>yersinia</code>	DoS tool for DHCP starvation attack
<code>dhcpstarv</code>	DoS tool for DHCP starvation attack
<code>dhcpiq</code>	DoS tool for DHCP starvation attack
<code>nbtscan</code>	Network scanner for NetBIOS name information
<code>p0f</code>	Passive traffic fingerprinting tool. Identifies hosts performing any incidental TCP/IP communication
<code>bannergrab</code>	Network service banner grabbing tool. Sends a trigger to the service and collects basic information
<code>nscan</code>	Fast network scanner optimized for Internet-wide scanning
<code>zmap</code>	Fast single packet network scanner. An improved version of <code>nmap</code> designed for Internet-wide scanning
<code>masscan</code>	Fast Internet port scanner
<code>fragrouter</code>	IDS evasion toolkit. Reroutes network traffic
<code>dnsspoof</code>	DNS spoofing tool. Forges replies to DNS queries on the LAN
<code>responder</code>	LLMNR, NBT-NS, and MDNS poisoner
<code>scapy</code>	Packet manipulation tool. Features packet forging, decoding, injection, and other network operations
<code>mitmf</code>	Framework for MitM attacks
<code>loki</code>	Firewall evasion tool that encapsulates commands into the payload of ICMP packets
<code>hts</code>	HTTPTunnel server. Used in conjunction with the HTTPTunnel client <code>htc</code> to tunnel network connections through pure HTTP traffic (GET and POST requests), hence bypassing restrictive firewalls or proxies
<code>htc</code>	HTTPTunnel client
<code>iodined</code>	Firewall evasion tool. Tunnels IPv4 traffic through a DNS server. Replaces the obsolete <code>tcp-over-dns</code>
<code>iodine</code>	Client for <code>iodined</code>
<code>loic</code>	Low Orbit Ion Cannon, a GUI tool for network stress testing and DoS/DDoS attacks
<code>hoic</code>	High Orbit Ion Cannon, a GUI tool for network stress testing and DoS/DDoS attacks
<code>zz</code>	Zombie Zapper, a countermeasure tool capable of stopping DDoS packet flooding attacks carried out by Trin00, TFN, Stacheldraht, etc.

kismet	Wireless sniffer and IDS
fern-wifi-cracker	Wi-Fi auditing and attack tool, with GUI
airsnort	WEP key cracker using the Fluhrer, Mantin, and Shamir attack (FMS)
wepattack	WEP key cracker. Takes a dumpfile as input
WEPCrack	WEP key cracker
airfart	GUI tool that detects Wi-Fi devices and displays their signal strength
cowpatty	WPA-PSK key cracker via dictionary attack. Part of the coWPAtty package
genpmk	WPA-PSK key cracker via precomputation attack
hackrf_info	Tool that probes the HackRF One device and shows its configuration. HackRF One is an open source SDR (Software Defined Radio) peripheral for testing RF systems. It is able to transmit and receive radio signals in the 1 MHz - 6 GHz range and can be used to carry out attacks such as replay, jamming, fuzzing, BlueBorne, etc. Usually used with the Ubuntu, Gentoo, and Pentoo distros
rfcat	ISM band radio multipurpose tool
rfcat-rolljam.py	Exploit for the rolling code attack. Jams, captures, and replays radio signals
scapy-radio	Modified version of the <code>scapy</code> packet manipulation tool to include RF capabilities
RFCrack.py	Radio Frequency communications scanner and test workbench. Can be used to perform the rolling code attack
btscanner	Bluetooth scanner with ncurses UI
bluescan	Bluetooth scanner
bluediving	Bluetooth pentesting suite
bluesniff	Bluetooth device discovery (wardriving) utility. Proof of concept

<code>sqlmap</code>	Automatic SQL injection exploitation/pentesting tool. The GUI version is Tyrant SQL
<code>mole</code>	Automatic SQL injection exploitation/pentesting tool
<code>jsql</code>	jSQL Injection, a GUI tool to retrieve database information from a remote server
<code>sqlsus</code>	Fast and efficient SQL injection and takeover tool for MySQL
<code>ISR-sqlget</code>	Blind SQL injection tool
<code>fatrat</code>	TheFatRat, a Remote Access Trojan and exploiting tool
<code>horsepill</code>	Ramdisk-based rootkit
<code>traitor</code>	Tool for privilege escalation. Uses several methods to exploit misconfigurations and vulnerabilities to obtain root shell access
<code>amap</code>	Application Mapper, a scanning and fingerprinting tool for pentesting
<code>amapcrap</code>	Fuzzer that sends random data to a TCP/UDP port and gathers a response for use by <code>amap</code>
<code>svmap</code>	Scanner for discovering VoIP SIP devices. Part of the SIPVicious package
<code>svwar</code>	Scanner for identifying active extensions on a PBX
<code>svcrack</code>	Online password cracker for SIP PBX
<code>svcrash</code>	Countermeasure for unauthorized <code>svwar</code> and <code>svcrack</code> scans
<code>svreport</code>	SIP audit session manager and reports exporter
<code>ike-scan</code>	Tool for discovering IKE hosts (IPsec VPN servers) and determine which IKE implementation they use
<code>ddosim</code>	Layer 7 DDoS simulator. Simulates zombie hosts connecting to a specific application on the target host
<code>fat</code>	Firmware Analysis Toolkit, a toolkit to find and analyze vulnerabilities in the firmware of IoT devices and embedded devices
<code>radamsa</code>	Mutation-based, general purpose fuzzer
<code>afl-fuzz</code>	American Fuzzy Lop, a coverage-guided fuzzer that uses a genetic algorithm
<code>balbuzard</code>	Tool to extract patterns of interest (e.g. URLs, IP addresses, typical malware strings) from malicious files
<code>bbcrack</code>	Tool to brute force typical malware obfuscation transforms (XOR, ROL, ADD, etc.) and discover the algorithms and keys used. Based on patterns of interest
<code>bbharvest</code>	Tool to extract all patterns of interest found when applying typical malware obfuscation transforms (XOR, ROL, ADD, etc.) trying all possible keys
<code>bbtrans</code>	Tool to apply any of the malware obfuscation transforms from <code>bbcrack</code> to a file
<code>binwalk</code>	Tool to search, identify, and extract files and executable code embedded inside firmware images. It is also able to calculate the entropy of a file to detect if it's a random, encrypted, or standard file
<code>pspy</code>	Process snooper which does not need root permission. It works by using inotify watchers, which trigger scans of procs

w3af	Web Application Attack and Audit Framework, GUI. Finds and exploits web application vulnerabilities
w3af-console	Web Application Attack and Audit Framework, CLI version
nikto	Webserver vulnerability scanner
skipfish	Webserver vulnerability scanner
uniscan	Simple webserver vulnerability scanner
burp	Burp Suite, a comprehensive testing platform for web applications security. Allows intercepting and modifying HTTP/HTTPS requests and replies to perform HTTP session hijacking
zap	OWASP Zed Attack Proxy, a pentesting GUI tool for web applications
webScarab	Tool for testing the security of web applications. Runs as a proxy
arachni_web	Modular framework and penetration testing platform to test the security of web applications
nuclei	Pentesting tool, based on templates, for configurable targeted scanning
xsser	Automatic framework to detect, exploit, and report XSS vulnerabilities in web applications
vega	GUI tool to test the security of web applications
paros	GUI tool with HTTP/HTTPS proxy functionality for assessing web application vulnerability
wapiti	Vulnerability scanner for web applications
httpbee	Web application testing tool
wsfuzzer	Fuzzer for black box testing of web services (HTTP and SOAP)
ffuf	Fast web fuzzer
httprint	Web server fingerprinting tool
wafw00f	Tool that detects and fingerprints a WAF (Web App Firewall) i.e. a firewall that protects a webserver
wafninja	Tool that circumvents a WAF by automating the steps necessary for bypassing input validation
whatweb	Web scanner. Detects and identifies web technologies, plugins, JavaScript libraries, etc.
sublist3r	Enumerator of website subdomains. Uses common web search engines
scrapy	Application framework for web scraping, web crawling, data mining, and other web content extraction
dirbuster	Web spider with GUI that attempts to find via brute force all non-linked (hidden) directories and files
sf	Spiderfoot, an OSINT website reconnaissance tool
dirsearch	Tool that brute forces directory and file names in web servers
dirbuster	Tool that brute forces directory and file names in web servers
gobuster	Tool that brute forces directory and file names in web servers, DNS subdomains, Virtual Host names, Open Amazon S3 buckets, Open Google Cloud buckets, and TFTP servers
feroxbuster	Tool for Forced Browsing
ferret	Cookie sniffer
hamster	Session hijacker. It runs as a proxy, reusing cookies stolen by <code>ferret</code> from other clients' sessions
fiddler	Web debugging proxy tool, with GUI. Can be used to test the security of web applications
sqlninja	Tool to exploit SQL injection vulnerabilities against web applications that use Microsoft SQL Server as database backend
SQLiX	Web crawler that attempts to find SQL injection vulnerabilities on explored websites
sslstrip	Tool for the HTTPS Stripping attack
slowloris	Tool for the DoS Slowloris attack against web servers
slowhttptest	Tool for testing slow HTTP DoS attacks such as Slowloris, Slow Read, R U Dead Yet, etc.
wpscan	WordPress vulnerability scanner

<code>john</code>	John The Ripper, a password cracker
<code>hashcat</code>	Password cracker and password recovery tool
<code>dsniff</code>	Password sniffer, able to operate over many cleartext network protocols
<code>hydra</code>	Login cracker tool supporting many network protocols and password cracking techniques
<code>medusa</code>	Brute force login cracker
<code>ncrack</code>	High-speed network authentication cracker
<code>ophcrack-cli</code>	Password cracker for MS Windows passwords. Uses rainbow tables
<code>patator</code>	Multi-purpose brute forcer (for FTP, SSH, Telnet, SMTP, DNS, MySQL, PostgreSQL, etc.)
<code>rcrack</code>	Hash cracker that uses rainbow tables. Part of the RainbowCrack package
<code>rtgen</code>	Rainbow table generator for <code>rcrack</code>
<code>crunch</code>	Wordlist generator. Creates a list of words based on permutation and combination of a character set

honeyd	Honeypot daemon. It creates virtual hosts, and simulates their networking stack and any desired network service
labrea	Honeypot for incoming IP connections. Replies to unanswered ARP requests, creating a virtual host with the related unused IP address, which then performs Layer 4 tarpitting
sshhipot	SSH honeypot
artillery	Honeypot with monitoring and alerting system
honeytrap	Extensible toolkit for running and monitoring honeypots
kojoney	Low-interaction honeypot that emulates an SSH server
honeypy	Medium-interaction honeypot
cowrie	High-interaction SSH and Telnet honeypot

snort	Network IDS/IPS with real-time traffic analysis and packet sniffing. It is configured via a customized ruleset that uses <code>libpcap</code>
ossec-hids	OSSEC, a HIDS with log monitoring and SIEM capabilities
tripwire	HIDS with integrity-based detection of unauthorized filesystem changes
aide	Advanced Intrusion Detection Environment, a HIDS with integrity-based detection. It makes a snapshot of the filesystem state and records it in a database, to check integrity of files at a later time
nessusd	Nessus, a well-known and comprehensive vulnerability scanner
saint	Security Administrator's Integrated Network Tool. Vulnerability scanner, originally based on SATAN (Security Administrator Tool for Analyzing Networks)
oscap	OpenSCAP command line tool. SCAP (Security Content Automation Protocol) is a framework of specifications for automated configuration, vulnerability and patch checking, and security compliance
lynis	Security auditing tool. Checks system and software configuration
debsecan	Debian Security Analyzer. Reports vulnerabilities in packages installed on the system
bastille	OS hardening tool. Obsolete, reported not working with Linux kernel 3.0 and later
rkhunter	Rootkit Hunter. Scanner for rootkits, backdoors, and other malware
chkrootkit	Check Rootkit. Scans the system for rootkits by identifying signatures of core system programs and detecting inconsistencies between the <code>/proc</code> filesystem and the <code>ps</code> command's output
pentbox	Security suite including password crackers, honeypots, DoS tools, etc.
websploit	Exploit framework containing reconnaissance and attack tools for various technologies
psad	Port Scan Attack Detector. Uses <code>iptables</code> log messages to detect and block port scans and other malicious network traffic
nexphisher	Automated phishing toolkit featuring many social media websites
stegdetect	Detector of steganographic content in graphic image files
inspy	LinkedIn enumerator. Attempts to find technologies and people at a specified target company
recon-ng	Web reconnaissance framework
dog	Recon Dog, an OSINT reconnaissance tool
maltego	OSINT tool with GUI that visualizes discovered data in a graph format for link analysis
JustMetadata	OSINT tool that gathers information about a large number of IP addresses and attempts to extrapolate relationships between them

Advanced persistent threat (APT)	Stealthy attack where the attacker gains unauthorized access to a system and remains undetected for a long period.
Zero-day attack (0day)	Attack exploiting a software vulnerability that is still unknown or for which no fix exists yet.
Man-in-the-middle (MitM)	Network-based threat where the attacker inserts itself undetected in the communication channel between two legitimate parties (network-level hijacking) and then proceeds to sniff, relay, and possibly modify the traffic. Countermeasure: mutual authentication of parties.
Replay attack Playback attack	Attack where the attacker eavesdrops on a communication, then maliciously sends again parts of a valid data transmission. Countermeasure: data tagging e.g. nonces, rolling code.
Rolling code attack Hopping code attack	Attack against the rolling code (used itself as a defense against replay attacks) used in keyless systems. The attacker jams the signal and sniffs a first code sent by the target. As the first code did not have any effect, the target sends a second code which is sniffed too by the attacker; at the same time the attacker forwards the first code which is received by the system, but the target believes it is the effect of the second code. Later on, the attacker uses the second code to gain unauthorized access to the system.
Side channel attack	Attack based on information obtained from the implementation of a system (e.g. analysis of power consumption, timing, electromagnetic leaks, sound) and not from weaknesses in the algorithm itself (e.g. cryptanalysis, software bugs).
Supply chain attack	Attack against the less secure elements in an organization's supply chain, usually done by tampering with the manufacturing process of the end-user software or hardware appliance (e.g. installing a backdoor in the firmware of a router). Countermeasure: use a SBOM (Software Bill Of Materials) to analyze vulnerabilities.
Slopsquatting	Supply chain attack consisting in listing non-existent software package names hallucinated by AI coding assistants, creating malicious software packages with such names, then waiting for an AI to re-hallucinate their names and embed the malware within their AI-generated code.
Warshipping	Attack consisting in using a physical package delivery service to deliver an attack vector (e.g. a backdoored router) to a target.
Banner grabbing	Reconnaissance technique consisting in initiating a connection to the desired service and noting the software type and version mentioned in the service banner. Countermeasure: configure banners to show minimal information.
Username enumeration	Reconnaissance technique in which the attacker tries to determine whether a specific username exists or not in the target system, or attempts to obtain a list of valid users. Countermeasure: configure the system to show minimal information about a failed login.
Google hacking Google dorking	Reconnaissance technique consisting in using advanced operators with specific strings (i.e. dorks) in the Google search engine to find specific versions of vulnerable web applications, misconfigurations, administration panels, sensitive files not supposed to be publicly accessible, etc.
Man-in-the-mobile	Infection of a mobile device with malware to bypass 2FA, as the malware relays the information to the attacker.
Privilege escalation	Host-based threat consisting in illegally gaining elevated access to resources that are normally protected from a program or user.
Confused deputy attack	Type of privilege escalation consisting in tricking a legitimate, more privileged program into misusing its authority on the system.
Sybil attack Pseudospoofing	Act of subverting a system by creating multiple fake identities. This may allow the attacker to e.g. acquire a disproportionate level of control over a reputation system by affecting voting results, or disrupt statistics about vehicular traffic.

Social engineering	Wide range of non-technical attacks consisting in deception and psychological manipulation of the target individual into divulging confidential information or performing unwarranted actions.
Pretexting	Social engineering attack where the attacker invents an elaborate scenario (i.e. a pretext) to engage the target individual under a fake identity and convince them to divulge confidential information or perform unwarranted actions.
Phishing	Social engineering attack aimed at obtaining sensitive information from people via a fake but legitimately-looking website controlled by the attacker. Usually perpetrated via an email message containing an obfuscated link to the malicious website.
Spear phishing	Personalized phishing targeted at a specific individual.
Whaling	Phishing targeted at a high-value individual (CEO, CISO, etc.).
Vishing	Phishing via VoIP.
Smishing	Phishing via SMS.
Quishing	Phishing via QR code. Is often carried out by covering a legitimate QR code in a public area (signs, parking meters, charging stations, advertisements, etc.) with a sticker containing a QR code to the malicious website.
Tailgating Piggybacking	Social engineering attack in which an attacker lacking proper authorization follows an authenticated individual into the targeted restricted area.
Shoulder surfing	Act of getting access to sensitive information by spying an individual entering the data.
USB drop attack Baiting	Social engineering attack consisting in leaving a bulk of malware-infected USB flash drives in public places for people to find and use.
Dumpster diving	Act of searching through discarded paper documents, left behind by the target organization, in order to find and exploit information.
Rubber hose cryptanalysis	Euphemism for extracting cryptographic secrets from the target by means of coercion or violence.
Black bag cryptanalysis	Euphemism for obtaining cryptographic secrets from the target by breaking and entering the premises, burglary, theft, etc. This term is also used for other non-cryptanalysis methods such as keystroke logging, infection via virus or trojan horse, etc.

Denial of Service (DoS)	Cyberattack towards hosts or networks, aimed at preventing or reducing availability of services to legitimate users. Countermeasures: blackhole filtering aka null routes, to drop all traffic coming from the attacker. Detection via activity profiling, sequential change-point detection (Cumulative Sum algorithm), wavelet-based signal analysis of traffic's spectral components.
Distributed Denial of Service (DDoS)	DoS launched simultaneously from several attacking hosts (usually a group of compromised machines i.e. a botnet).
Distributed Reflected Denial of Service (DRDoS)	DDoS carried out by forging requests to a large number of remote hosts using the target host's spoofed source IP address.
Permanent Denial of Service (PDoS) Phlashing	Hardware-targeted DoS which replaces the target device's firmware with a faulty one, bricking the device permanently.
Multi-vector attack	DoS combining volumetric, protocol, and application-layer attacks.
Regular expression Denial of Service (ReDoS)	DoS in which the attacker supplies maliciously crafted input to a program's regular expression, causing it to hang for an excessive amount of time. This occurs because many regex engine implementations can encounter extreme cases (e.g. backtracking) that lead to very slow execution.

Ping of death	A malformed or oversized ping packet which, when reassembled by the target host, causes a buffer overflow, crashing or infecting with malicious code the target host. Obsolete; modern TCP/IP stacks, firewalls, and IDSs easily identify and discard pings of death.
Ping flood ICMP flood	DoS in which the attacker sends a large number of ICMP Echo Request packets to the target host.
Smurf attack	DRDoS in which the attacker sends a large number of ICMP Echo Request packets to a network broadcast address using the target host's spoofed source IP address. Obsolete; by default, ICMP requests to broadcast addresses are not forwarded anymore by routers.
Fraggle attack	DRDoS in which the attacker sends a large amount of UDP traffic to ports 7 (Echo Protocol) and 19 (CHARGEN) of multiple remote hosts, using the target host's spoofed source IP address.
SYN flood	DoS in which the attacker sends a large number of TCP SYN packets to the target host but never responds to its SYN/ACK, hence never completing the TCP three-way handshake. This creates a large number of half-open connections which, until they time out, fill up the target host's connection queue, preventing legitimate clients to connect to it. Countermeasures: TCP Intercept (aka SYN Proxy or SYN cookies).
Teardrop	DoS in which the attacker sends mangled IP fragments with overlapping and oversized payloads to the target host, causing it to crash. Obsolete; fixed in the Linux kernel v2.1.63.
Fragmentation attack	DoS performed by sending a large number of fragmented TCP or UDP packets to the target host, which will consume resources in reassembling and inspecting them.
Shrew attack	Low-rate DoS that exploits the retransmission timeout (RTO) mechanism of TCP. Performed by sending out a burst of traffic to a bottleneck router at the same time the client sends a request to the server. Hence, the router suspends the data transmission, packets are dropped during the RTO, and after the RTO the client needs to resend the lost packets, slowing down the transmission.
MAC spoofing	Act of setting a spoofed MAC address in the NIC in order to divert communications to the host controlled by the attacker. Countermeasure: IP Source Guard (on switches) using the DHCP snooping binding table.
MAC flooding	Dispatch of multiple Ethernet frames with different source MAC addresses to a switch. This fills up a switch's CAM table and forces the switch to failover to hub mode (i.e. broadcasting to all switch ports), allowing the attacker to sniff all network traffic in the LAN. Countermeasure: port security with max one MAC address per interface (on switches).
ARP spoofing ARP cache poisoning ARP poisoning ARP poison routing	Dispatch of forged ARP messages into the LAN to associate the attacker's MAC address with another host's (often the gateway) IP address, diverting communications to the attacker. Often the first choice of attack for the purpose of sniffing, or in preparation to MitM attacks or session hijacking attacks. Countermeasures: Dynamic ARP Inspection (on switches) using the DHCP snooping binding table, disabling of gratuitous ARP, static IP addresses and ARP tables. On wireless networks, Client Isolation (on the wireless router) which prevents wireless clients from communicating between them.
Port stealing	ARP spoofing aimed to associate the attacker's MAC address with another host's IP address on a switch's CAM table, which will then forward packets through the wrong switch port.
IP spoofing	Act of setting a spoofed IP address in the NIC so the attacker's host can appear to be some other host. Countermeasure: direct TTL probes (however, this works only if the attacker's host is in a different subnet).

DNS spoofing	Tampering with the name resolution mechanism of the target host so that a domain name resolves to an incorrect IP address controlled by the attacker. Can be carried out either by DNS hijacking, by deceiving the target host to use a rogue DNS server, or by tampering with the hosts file of the target host. Can be done in preparation to a MitM attack, although for a LAN an ARP spoofing would serve the same purpose and is easier to do. Countermeasures: DNSSEC, restriction of DNS service, master-slave DNS setup with no Internet access for the master, DNS anti-spoofing.
DNS hijacking	Attack which consists in compromising a DNS server and changing the mapping settings to redirect towards a rogue DNS server. This can also be done by stealing the domain name upon the sponsoring domain name registrar accredited by the ICANN (which manages the DNS root zone). Countermeasure: at the registrar level, REGISTRAR-LOCK status code to prevent unauthorized changes to the domain name.
Cybersquatting	Registration of a domain name which is similar to a well-known domain, product, or entity, in order to deceive users. Can be done in preparation to phishing attacks or scams.
Typosquatting	Cybersquatting where the attacker relies on typos or spelling mistakes made by users when they manually type a URL into a web browser.
Punycode attack IDN homograph attack	Cybersquatting where the malicious domain name is visually identical to a well-known domain name by switching one or more letters in a different alphabet, e.g. by replacing the Latin character "a" (Unicode U+0061) with the Cyrillic character "a" (U+0430).
Domain sniping Domain snapping	Registration of a domain name that has just expired, with the purpose of reselling it to the original owner at a higher price.
DNS cache poisoning	Injection of forged DNS records in the DNS resolver's cache, causing the name server to return an incorrect IP address for a domain name, hence redirecting traffic to the attacker.
Blind response forgery	DNS cache poisoning attack carried out by guessing the transaction ID (birthday paradox). Countermeasure: randomization of UDP source port.
DNS water torture	DDoS done by performing a large number of DNS queries for nonexistent subdomains of a target domain. Subdomains strings are randomly-generated by the attacker, hence the queries bypass the DNS cache and hit the DNS Authoritative Servers of the target domain.
DNS amplification attack	DRDoS in which the attacker sends a large amount of DNS queries to the target host's DNS server, using the target host's spoofed source IP address. The recursive resolution of queries ends up overwhelming the target host's DNS server.
DNS sinkhole attack Blackhole DNS attack	Act of providing incorrect DNS information to systems so to redirect their communications to a single destination. This can also be done for beneficial purposes, e.g. to block ads or stop botnets from contacting their C&C (Command and Control) host.

DHCP spoofing	Attack consisting in setting up a rogue DHCP server and use it to send forged DHCP responses to hosts. Often done to replace the IP addresses of the default gateway and DNS server, redirecting traffic to attacker-controlled nodes. Countermeasure: DHCP snooping and Dynamic ARP Inspection (DAI) on routers.
DHCP starvation	DoS in which the attacker floods a DHCP server with DHCP requests from spoofed MAC addresses, depleting the server's IP address pool and making it unable to allocate them for legitimate clients. Also done in preparation to the deployment of a rogue DHCP server.
IRDP spoofing	Injection of forged IRDP Router Advertisements to add default route entries to a target host, redirecting traffic to the attacker-controlled node.
Sinkhole attack	Attempt to attract network traffic by advertising fake routing updates. Once traffic passes through the malicious node, the attacker may alter the payload, launch a blackhole or wormhole attack, or perform other disruptive activities.
Blackhole attack Packet drop attack	DoS attack where an attacker-controlled node discards packets instead of relaying them. This can be done partially and/or selectively (e.g. depending on the time of the day, the source, the destination) in order to avoid detection.
Wormhole attack	Attack (usually carried out on wireless networks) where the attacker records packets in one location and then tunnels them to another location, selectively or as a whole.
Man-in-the-Cloud (MitC)	MitM-like attack against cloud file synchronization services, carried out by stealing and reusing a synchronization token from the target cloud user to obtain access to their files. Countermeasure: hardened policies for token expiration.
Wardialing	Reconnaissance technique consisting in automatically dialing every telephone number from a list (usually in a local area code) searching for modems, BBS, or fax machines. Obsolete, as dial-up Internet connections have mostly disappeared.
Efail	Attack in which the attacker is able to read the decrypted content of an email encrypted via OpenPGP or S/MIME, if the mailclient is allowed to load remote content or run active content e.g. JavaScript. Countermeasure: a correctly configured and up-to-date mailclient.
Blast-RADIUS	Attack against the RADIUS protocol affecting RADIUS implementations which use non-EAP authentication methods over UDP. It consists in a MitM between RADIUS client and server to forge a valid Access-Accept message in response to a failed Access-Request, hence giving the attacker access to protected resources.
Website fingerprinting	Traffic analysis technique in which a local passive eavesdropper is able to determine the web pages visited by a client which uses privacy technologies (e.g. proxies, VPNs, or Tor) by applying machine learning algorithms to the client's network packets.
DefecTor	DNS-Enhanced Fingerprinting and Egress Correlation on Tor. Enhanced website fingerprinting attack done via the observation of DNS traffic from Tor exit relays.
Flow correlation	Traffic analysis technique in which an attacker passively monitoring the ingress and egress of an anonymity system is able to correlate traffic and de-anonymize clients by using side-channel information (e.g. network packets timings and sizes).
Browser-based attack [Tor]	Traffic analysis technique aimed at de-anonymize Tor clients, in which an attacker uses a malicious Tor exit relay to insert HTTP content (either an invisible iframe and a unique cookie, or some JavaScript code) referring to a malicious web server. The client's web browser will then make a direct connection to the web server.
Browser-based timing attack [Tor]	Traffic analysis technique aimed at de-anonymize Tor clients, in which an attacker uses a malicious Tor exit relay to insert HTTP content (a HTML meta refresh tag) referring to a malicious web server. The attacker generates the desired traffic signal by dynamically varying the webpage refresh delays.

Evil twin attack	Attack consisting in setting up a legitimately-looking rogue Wi-Fi Access Point to lure clients into connecting to it and then perform eavesdropping or MitM attacks. To improve effectiveness, the rogue AP can even transmit with a stronger signal. Countermeasures: network management software (on the network management side) with wired side inputs to detect devices connected to the LAN and hence also rogue APs; WIPS.
KARMA attack	Variant of the evil twin attack. Some vulnerable devices broadcast the list of their preferred networks (i.e. the SSIDs of APs to which they have already connected and are going to connect automatically). Upon receiving this information, an attacker can set up a rogue AP with a SSID from the list.
Client misassociation	Attack similar to the evil twin attack, consisting in setting up a rogue AP that duplicates the SSID and the MAC address of a legitimate AP.
Disassociation attack Deauthentication attack	Availability attack carried out by sending deauthentication frames to the AP to disconnect clients. This attack can be done against a specific client (by using the target client's spoofed MAC address) or all clients.
Beacon flood attack	Availability attack carried out by sending a large number of forged Wi-Fi beacons to confuse wireless clients and make it harder for them to connect to a legitimate AP.
Clear channel assessment attack Queensland attack	Physical layer DoS attack that exploits the CSMA/CA Clear Channel Assessment (CCA) to make the channel appear busy.
Fluhrer, Mantin, and Shamir attack (FMS)	Attack which exploits a weakness in the RC4 key scheduling algorithm to reconstruct the key from encrypted messages. This attack can be used to recover a WEP key.
Chopchop attack	Attack carried out against a WEP-encrypted wireless communication which allows to recover the unencrypted payload. The attacker chops off the last byte of data from a WEP-encrypted packet, replaces that byte, recalculates the checksum, and sends the packet to the AP. The AP discards the packet, until by trial and error the attacker eventually replaces a valid checksum and the AP accepts it. The same attack can be carried out against WPA-TKIP. This attack does not recover the WEP key.
Key Reinstallation Attack (KRACK)	Attack against the four-way handshake in the WPA2 authentication protocol. The attacker captures and replays the message in step 3 (containing the AP's nonce) to force nonce reuse; this allows the attacker to decrypt all traffic. Countermeasures: update all wireless devices with the latest security patches, patch the AP's firmware, use HTTPS, enable 2FA.
Fragmentation attack [WEP]	Attack consisting in extracting some keying information from a WEP packet, then sending ARP and LLC packets to the AP which resends them back, then extracting more keying information from the packets. This cycle is repeated until the attacker gets 1500 bytes of the PRGA (Pseudo Random Generator Algorithm) which can then be used to forge and inject packets. This attack does not recover the WEP key.
Caffe Latte attack	Attack allowing to recover a WEP key from a client by capturing an ARP packet from the client, manipulating it, and sending it back to the client.
Hirte attack Client-oriented fragmentation attack	Extension to the Caffe Latte attack, performed using any ARP or IP packet.
Wardriving	Detection and reconnaissance of WLANs by listening to SSID broadcasts or by sending probe requests, usually done from a moving vehicle.
Warchalking	Technique of advertising discovered WLANs in range by drawing specific symbols with chalk in public places, usually on pavements or walls.

Bluejacking	Sending of anonymous messages (e.g. spam) to a Bluetooth device, done by inserting the message in the BT connection request. Uses the OBEX (Object Exchange) protocol.
Bluesnarfing	Theft of information from a Bluetooth device. The attacker connects to the target BT device and performs a GET operation for known or guessed filenames. Carried out by exploiting a vulnerability in the OBEX protocol.
Bluebugging	Unauthorized remote access and takeover of a Bluetooth device.
Blueprinting	Footprinting performed against a Bluetooth device.
Bluesmacking	Ping of death attack carried against a Bluetooth device.
BlueBorne	Vulnerability in the Bluetooth implementation on multiple OSes that allows an attacker to take control of the target device, even if the device is not paired or even set to discoverable mode. The attacker gets the MAC address and performs OS fingerprinting on the device, then uses a BT exploit.

Network-level hijacking	Interception of TCP or UDP packets during transmission between client and server. This term is also used for the takeover of a legitimate TCP communication between two hosts, done via IP spoofing and MitM, sometimes using source routed packets. The attacker sniffs (or tries to predict) TCP Sequence and Acknowledgment numbers from the client, sends forged TCP Sequence and Acknowledgment numbers to the server to desynchronize the client, and finally inserts itself in the TCP session.
Application-level hijacking Session hijacking	Takeover of an HTTP session, usually done by stealing an HTTP session token. This is not an instance of network-level hijacking.
TCP/IP hijacking	Network-level hijacking in which the attacker sniffs the communications between two hosts to get the target host's Initial Sequence Number (ISN). The attacker then sends a packet with the target host's spoofed source IP address using the captured ISN. The other host receives the packet, increments its TCP Sequence number, and sends an ACK to the target host which ignores it. The attacker continues to send spoofed packets with forged TCP Sequence and Acknowledgment numbers, causing the target host to have desynchronized values and making its connection hang. At this point, the attacker inserts itself in the TCP session, replacing the target host.
Blind hijacking	Network-level session hijacking in which the attacker tries to predict ISN and TCP Sequence and Acknowledgment numbers, without being able to see the response. Can be used to inject malicious data into the communication, and does not require source routing. This is not a MitM attack.
RST hijacking TCP reset attack	Injection of an RST packet with spoofed source IP address within a legitimate TCP communication, to terminate the connection. May be done in preparation to TCP/IP hijacking.
UDP hijacking	Network-level hijacking where the attacker forges UDP replies from the server.
Source routing attack	Network-level session hijacking that uses the source routing field in the IP header to specify a packet route so to, with the help of a trusted host, divert packets towards the attacker's node. Used in IP spoofing attacks. Obsolete; by default, network devices nowadays discard source routed packets.
Session fixation	Application-level session hijacking in which the attacker sets a session ID on behalf on the target host. This can be done e.g. via a phishing email. This attack is effective against e.g. web applications that do not change the session cookie after a successful login and instead allow additional privileges to it.
Session prediction	Application-level session hijacking in which the attacker predicts a session ID value. The attacker needs beforehand to collect valid session ID values that identify authenticated users, and to analyze and understand the session ID generation algorithm.
Session brute-forcing	Application-level session hijacking in which the attacker tries all possible session ID values until they successfully get access to the application.
Session riding	Application-level session hijacking obtained via Cross-Site Request Forgery. Countermeasures: check the HTTP Referrer header, ignore URL parameters when processing an HTTP POST command.
Session sidejacking Sidejacking Cookie stealing	Application-level session hijacking in which the attacker sniffs a session ID (session cookie) from a legitimate session and then reuses it to impersonate the legitimate client.

HTTP response-splitting attack HTTP header injection	Attack which consists in adding header response data into an input field so that the webserver splits the HTTP response into two; these can either be served to the attacker, served to a legitimate client, or discarded. Countermeasure: validation of client input. In particular, CR (%0d, \r) and LF (%0a, \n) characters should never be allowed in input.
HTTP request tampering	Unauthorized access to a web application obtained either by tampering with the URL query string, or by modifying the HTTP headers sent to the webserver e.g. the <code>Referer</code> : header which in vulnerable applications is used for access control.
HTTP Parameter Pollution (HPP)	Evasion technique used to bypass WAF security filters which consists in crafting an HTTP request containing multiple instances of a parameter with the same name, hence splitting the attack vector.
HTTP Parameter Fragmentation (HPF)	Evasion technique, often used along with HPP, which allows to reconstruct the parameter string passed in the HTTP request.
Webcache poisoning	Attack where the attacker uses a specially crafted request to force the webserver's cache to flush its contents and insert a URL with infected content in the cache, which is then served to legitimate clients accessing the cache. Can be performed via an HTTP response-splitting attack.
Directory traversal	Unauthorized access to directories outside the webserver's root directory, done by using repeatedly the <code>../</code> sequence in URLs.
Forced browsing	Attack aimed at enumerating and accessing resources not referenced (linked) by the web application, but nonetheless accessible.
Unvalidated redirect	Phishing in which the URL is that of a legitimate site but contains a redirect to the malicious site.
Unvalidated forwarding	Unauthorized access to a restricted webpage obtained fraudulently via an embedded forward query on the URL.
CRIME	Compression Ratio Info-leak Made Easy. Exploit against authentication web cookies transmitted over compressed HTTPS and SPDY connections, which results in session hijacking.
BREACH	Browser Reconnaissance and Exfiltration via Adaptive Compression of Hypertext. Session hijacking exploit analogous to CRIME, but performed against HTTPS when using HTTP compression.
HTTPS Stripping SSL Stripping	MitM attack which consists in hijacking the connection establishment from the target host to a remote secure webserver, then transparently downgrading all HTTPS traffic to HTTP. This attack works for webserver with redirection from HTTP to HTTPS, and only if the initial request to the webserver is sent as HTTP; then, if the client does not explicitly specify HTTPS for links, the attacker intercepts all HTTP 302 redirections and sends the client the requested content through HTTP. Countermeasure: HSTS (however, browsers not accepting HSTS cookies will still be vulnerable).
Forbidden attack	Attack exploiting vulnerable implementations of the TLS protocol that incorrectly reuse the same cryptographic nonce when data is encrypted. This allows the attacker to sniff the HTTPS connection and inject content.
HTTP flood	Layer-7 DDoS targeted at webserver. Carried out by sending a large number of HTTP GET and POST requests towards the target webserver.
Slow HTTP attack	Low-bandwidth DoS targeted at webserver. Carried out by keeping several connections to the target webserver open as long as possible, depleting the webserver's connection pool. This is done by exchanging HTTP traffic at an extremely slow rate (1 byte/min or less), preventing the webserver from going into timeout for idle connection.
Slow Post attack	Slow HTTP attack carried out by sending correct HTTP requests and headers at an extremely slow rate.
Slow Read attack	Slow HTTP attack carried out by reading the webserver's response at an extremely slow rate.
Slowloris	Similar to the Slow Post attack, but carried out by sending partial HTTP request and headers, while never completing the request.
R U Dead Yet (RUDY)	Slow HTTP attack carried out by opening concurrent POST HTTP connections and delaying sending the body of the POST request.

Cross-Site Scripting (XSS)	Injection of malicious client-side scripts (e.g. via input textfields) in webpages, which are then executed when the webpage is viewed by other clients. This can lead to session hijacking, data theft, exploiting user privileges, etc. Identification of entry points for user input is the first step for exploiting a website. Countermeasures: set the HttpOnly flag in session cookies, sanitize user input.
Non-persistent XSS Reflected XSS	XSS attack where the malicious data provided by the attacker is used immediately by server-side scripts to display results to the targeted user only.
Persistent XSS Stored XSS	XSS attack where the malicious data provided by the attacker is saved by the server and permanently displayed to all users visiting the website.
Server-side XSS	XSS attack where the malicious data provided by the attacker is wholly processed server-side. Historically, this was the first kind of XSS attack.
DOM-based XSS	XSS attack where the malicious data provided by the attacker does not affect the webserver but it is reflected fully client-side, where all the presentation logic (often JavaScript) occurs.
Cross-Site Flashing (XSF)	XSS-like attack where the malicious data provided by the attacker is used in some specific video playing functions and variables by Flash scripts.
Cross-Site Tracing (XST)	XSS attack, involving the use of the HTTP TRACE method, that allows stealing cookies via Javascript.
Cross-Site Request Forgery (CSRF) One-click attack	Attack which consists in having the target client unknowingly submit a maliciously crafted web request to a webpage. Can be performed via specially designed image tags, hidden forms, JavaScript functions, etc. Countermeasure: random tokens in the web application.
Server-Side Request Forgery (SSRF)	Attack which consists in inducing the server-side application to make HTTP requests to an arbitrary domain. This results in unauthorized actions and access to data, and may lead to arbitrary command execution.
URL parameter tampering Web parameter tampering	Modification of parameters in the URL to exploit vulnerable applications that use them, e.g. <code>http://www.bank.com/account?id=345&amount=200000</code>
Session poisoning	Unauthorized access to a web application by exploiting (e.g. via URL parameter tampering) weak user input validation. The same term may also indicate the takeover of a session of a legitimate user e.g. by injecting malicious content, or via cookie poisoning.
Cookie sniffing	Unauthorized access to a web application by sniffing a cookie belonging to a legitimate user and reusing it to bypass the authentication process, ending up logged in as that user. Countermeasure: SSL, set the secure attribute on cookies (which will therefore be sent only over HTTPS).
Cookie replay	Replay attack in which the attacker captures a cookie belonging to a legitimate logged in user, then reuses it to perform malicious activities on behalf of that user. The attack persists until the user logs off.
Cookie poisoning	Unauthorized access to a web application by crafting a cookie, or by sniffing and modifying a cookie belonging to a legitimate user. Countermeasures: cookie expiration, associate cookie's credentials to an IP address.
Cookie parameter tampering	Unauthorized access to a web application by tampering with the parameters of a cookie and resubmitting it.
SSI injection	Code injection technique consisting in injecting scripts in webpages via SSI. SSIs (Server Side Includes) are directives present on web applications which allow inserting dynamic content into an HTML page before it is loaded or visualized.
CORS attack	Attack exploiting the Cross-Origin Resource Sharing (CORS) mechanism. CORS allows restricted resources on a webpage to be requested from another domain outside the domain from which the first resource was served; this is safer than allowing all cross-domain requests. It bypasses the Same-Origin Policy (SOP) which forbids certain cross-domain requests (e.g. Ajax).
Connection String Parameter Pollution (CSPP)	Injection of connection string parameters into other existing parameters, often carried out by using a ; character. This can be done e.g. in connection strings for backend databases.

Website defacement	Unauthorized changes made to the website appearance and content, to show the attacker's propaganda and/or infect legitimate clients with malware.
Pharming	Attack intended to redirect legitimate traffic from a website to a fake one. Usually carried out via DNS spoofing. Considered an advanced form of phishing.
Watering hole attack	Attack in which the attacker identifies which websites the target users visit more often and infects those websites with malware (e.g. via XSS).
Framing attack	Insertion of a malicious webpage inside a legitimate webpage by using the <code><iframe></code> (inline frame) HTML tag.
Clickjacking UI redress attack UI redressing (UIR)	Deceptive technique consisting in tricking web users to click on a different element from the one they think they are clicking. Usually done via a framing attack, by having an invisible iframe with malicious content on top of a visible iframe with innocuous content.
Man-in-the-browser (MitB)	Attack related to MitM where a Trojan horse infects a web browser in the target host, and injects HTML code in the browser's requests and responses. The Trojan operates between browser and OS API, allowing it to read data before encryption when it is sent from the host, and read data after decryption when it is received by the host.
Password spraying attack	Brute force attack consisting in trying different logins based on a list of usernames with the same default password. This is done to avoid account lockouts that would normally occur when trying different passwords for the same login. Countermeasure: implement brute force lockout on both username and password fields.
ClearFake ClickFix	Attack in which a malicious webpage displays a fake CAPTCHA or Cloudflare Turnstile notification requiring the user to paste the clipboard contents to a CLI prompt and run it. The clipboard contents consist of commands (usually obfuscated) that download a malicious payload and execute it on the user's machine. Countermeasure: avoid running commands of dubious origin.
In-session phishing	Phishing attack exploiting the ability of a malicious website to detect whether a user is simultaneously in a session on a legitimate website (e.g. an online banking platform) in the same browser. The attacker then opens a pop-up window that pretends to belong to the legitimate website, and which is used to collect user data. Obsolete, as it depended on a JavaScript handling vulnerability in browsers that has since been patched.

XML External Entity (XXE)	SSRF attack consisting in passing malicious XML input, referencing an external entity containing infected data, to an application with a weakly configured XML parser. This may cause confidential data disclosure, DoS, remote code execution, etc. XML DoS issues are a common SOA vulnerability.
XML bomb Billion laughs attack	Exponential entity expansion attack consisting in a specially crafted XML document where the root element contains a defined entity which contains several defined entities, each of which contains several defined entities, and so on. This takes up a large amount of memory and causes a DoS when the XML document is processed. Countermeasures: capping the memory allocated for an individual XML parser, treating entities symbolically and expanding them only when needed.
XML injection XML poisoning	Insertion of specially crafted XML fields in SOAP requests for web services. This may cause confidential data disclosure, DoS, remote code execution, etc.
XPath injection	Code injection technique in which the attacker provides malicious parameters to construct the wrong XPath query and access the wrong XML node. XPath is a query language for selecting nodes from an XML document.
Wrapping attack [XML]	Attack performed during the translation of a SOAP message in the TLS layer, by intercepting the message, adding the body to the header, injecting a malicious payload in the body, and sending the message to the server as a legitimate client. The server therefore verifies the XML signature as valid. Countermeasures: XML schema validation, authenticated encryption in the XML-Enc (XML Encryption) specification.

SQL injection (SQLI)	Code injection technique in which the attacker inserts malicious SQL statements into an input field for execution. This is done by using single quotes (') and double quotes ("). Countermeasures: user input sanitization, avoiding constructing dynamic SQL with concatenated user input values, keeping untrusted data separate from commands and queries, using least privilege account types for connection to the database.
In-band SQL injection	SQL injection where the attacker uses the same communication channel to send the query and retrieve the result. Most common type of SQL injection. Can be: Error-based: the attacker causes the database to throw an error, e.g. by forcing a conversion type, and gains knowledge by analyzing the error message Illegal query: the attacker sends a logically incorrect query and gains knowledge by analyzing the error message Union: the attacker uses the UNION operator to get the field values from other tables; can be coupled to other operators e.g. ORDER BY to find the number of fields in a table Tautology: the attacker inserts the OR operator with a tautology so that a WHERE clause is always true, e.g. ' OR '1'='1' End-of-line comment: the attacker uses -- to insert an end-of-line comment to make the database parser ignore the rest of the query Inline comment: the attacker uses /* */ to insert an inline comment to bypass fields Piggybacked query aka stacked queries: the attacker uses a ; character to insert additional malicious queries to the original query Stored procedure: the attacker input is fed to dynamic SQL statements Second order: the attacker's input is saved in a database and used later when the attacker submits a second query; requires knowledge of the application backend
Out-of-band SQL injection	SQL injection where the attacker retrieves the result via a different channel e.g. email, HTTP, or file I/O functions. More difficult to perform.
Blind SQL injection Inferential SQL injection	SQL injection attack used when the application does not show a useful error message and hence the attacker needs to assess whether the application might be vulnerable to an SQL injection. This attack is time-intensive. Can be: Time-based aka double blind: the attacker sends a query which may be true or false embedded with sleep or benchmark functions, and gets information from the time delay that occurs in the response Boolean-based: the attacker sends a query which causes the application to return a different result depending on whether the query returns true or false Heavy query: the attacker sends a query that takes noticeable time to execute, e.g. a <code>SELECT COUNT(*)</code> from multiple tables
ORM injection	SQL injection attack against a data access object model generated by ORM (Object Relational Mapping).

Buffer overflow Buffer overrun	Attack carried out by writing data to a buffer over the buffer's boundaries, overwriting the adjacent memory addresses. This allows the attacker to modify the target process' address space to control process execution, crash the process, and modify variables. Countermeasure: buffer bounds checking.
Fork bomb	DoS attack consisting in a process that recursively replicates itself ad infinitum, depleting system resources and slowing down or crashing the system due to resource starvation.
Fuzzing	Attack (also a software testing technique) consisting in sending invalid, unexpected, or random data as input to a program in order to crash it or provoke an exception e.g. a buffer overflow or a memory leak. A memory leak is a form of memory consumption where the program fails to release an allocated block of memory when it is no longer needed.
Shrink-wrap code attack	Attack consisting in exploiting holes in unpatched or misconfigured software (e.g. software with default insecure configuration options).
Code injection	Attack in which the attacker inserts text in a data field that gets interpreted as code.
File injection	Code injection technique in which the attacker exploits dynamic file include mechanisms in vulnerable web applications (especially in PHP). The attacker provides a URL pointing to the malicious file, which is used by the web application instead of the intended local file.
DLL injection	Attack consisting in forcing a process to load a dynamic-link library, resulting in the attacker's malicious code running within the address space of that process.
LDAP injection	Code injection technique in which the attacker inserts malicious LDAP user parameters into an input field to get access to the LDAP database. Countermeasure: user input sanitization.
Insecure deserialization attack	Injection of malicious code into a serialized object; if the application uses a vulnerable algorithm for deserialization, the code is executed when the object is deserialized. Serialization is the act of converting an object into a format (e.g. XML, JSON, YAML) which can be written to disk, transmitted over the network, or sent to a stream (e.g. stdout).
Cross-guest VM breach	Side channel attack carried out by running a VM on the same physical host as the target VM and taking advantage of shared physical resources (e.g. CPU cache) to extract cryptographic secrets.
Cryptojacking	Attack consisting in exploiting a device's processing power (CPU/GPU) to mine cryptocurrency without the owner's consent. This results in increased power consumption, degraded system performances, and potential hardware damage from overheating.
Shellshock	Vulnerability in the Bash shell which allows an attacker to execute arbitrary commands by exploiting the function export feature of Bash.
Log4Shell	Remote Code Execution vulnerability in the Apache Log4j library. An attacker causes the application to write one string into the log, then exploits the message lookup substitution function to upload malicious code via JNDI into the application.
Dirty Pipe	Vulnerability in the Linux kernel that allows an attacker to overwrite read-only files with arbitrary data, leading to privilege escalation and root access. Fixed in Linux 5.16.11.
Logjam	Vulnerability affecting Diffie-Hellman key exchanges with key size from 512 to 1024 bits, allowing an attacker to downgrade a TLS connection and read and modify the exchanged data.
Heartbleed	Buffer over-read vulnerability in the OpenSSL cryptography library. The attacker client sends a heartbeat request specifying a length larger than the message payload. The server replies filling up all the allocated buffer with the contents of the active memory, which may include private data.
POODLE	Padding Oracle On Downgraded Legacy Encryption. Vulnerability in the SSL v3.0 protocol that allows an attacker to eavesdrop on encrypted communications.

Evil maid attack	Attack carried out on an unattended device which is tampered by the attacker in an undetectable way (e.g. by replacing the firmware). Countermeasure: ensure permanent physical security of the device.
Cold boot attack	Attack relying on the data remanence property of RAM, whose stored contents can be read for seconds or minutes (or even days, if RAM chips are frozen with liquid nitrogen) following a poweroff. An attacker with physical access to the machine does a hard reset, then boots up the machine on a specialized live external OS (via USB drive, CD-ROM, etc.) and performs a memory dump of the machine. Typically used to retrieve encryption keys from a running OS.
Juice jacking	Attack consisting in injecting malware and/or surreptitiously copy sensitive data from a device (usually a smartphone or tablet) when the latter is connected to a USB port for charging. Countermeasure: deny USB data transfer from device, use a USB data blocker aka USB condom, use a trusted charger connected to the power line.
ChoiceJacking	Juice jacking attack in which the malicious charging station enables on its own the USB data transfer mode on the target device. Countermeasure: use a USB data blocker, use a trusted charger connected to the power line.
DMA attack	Attack consisting in accessing memory locations of a computer via high-speed expansion ports (USB, FireWire, Thunderbolt, PCI Express, etc.) that permit Direct Memory Access. Countermeasure: ensure physical security of the device, disable DMA capability on expansion ports when console is locked.
TEMPEST	Telecommunications Electronics Materials Protected from Emanating Spurious Transmissions. Side-channel attack consisting in eavesdropping data from a machine through leaking emanations such as radio signals, electrical signals, sounds, and vibrations.
Van Eck phreaking	TEMPEST attack exploiting the electromagnetic emissions of electronic devices, most commonly CRT displays.
Spectre	Vulnerability affecting the branch prediction mechanism (done to improve performances) on microprocessors. An attacker can exploit the side effects of speculative execution, and particularly branch misprediction, to access private data.
Meltdown	Vulnerability affecting some Intel x86, IBM POWER, and ARM microprocessors. An attacker can have a process, even without proper authorization, read any memory location by exploiting a race condition between memory access and privilege check during instruction processing.
Rowhammer	Attack consisting in accessing physical rows inside vulnerable memory chips millions of times per second, causing bit flips in neighboring rows. This allows the attacker to e.g. bypass security sandboxes and escalate privileges of untrusted applications.

Linear cryptanalysis	Cryptanalysis based on finding affine approximations to the actions of a cipher, analyzing pairs of plaintext and the corresponding ciphertext to try to recover the encryption key.
Differential cryptanalysis	Cryptanalysis based on the analysis of how differences in the input affect the output.
Integral cryptanalysis	Cryptanalysis based on the analysis of pairs of inputs differing in only one bit.
Known plaintext attack	Linear cryptanalysis technique where the attacker has access to some plaintext as well as the corresponding ciphertext.
Chosen plaintext attack	Cryptanalysis technique where the attacker is able to obtain the ciphertext corresponding to a plaintext of their choice.
Chosen ciphertext attack	Cryptanalysis technique where the attacker is able to obtain the plaintext corresponding to a ciphertext of their choice.
Adaptive chosen plaintext attack	Cryptanalysis technique where the attacker has access to the encryption device and is able to obtain the ciphertexts corresponding to plaintexts of their choice, making adaptive changes in the plaintext where needed.
Adaptive chosen ciphertext attack	Cryptanalysis technique where the attacker has access to the encryption device and is able to obtain the plaintexts corresponding to ciphertexts of their choice, making adaptive changes in the ciphertext where needed.
Non-adaptive chosen ciphertext attack Lunchtime attack	Chosen ciphertext attack where the attacker has access to a limited set of plaintexts and ciphertexts, or has access to the device for a limited time.
Related key attack	Cryptanalysis technique where the attacker is able to obtain the ciphertext corresponding to a plaintext of their choice, encrypted with two different but similar keys.
Chosen key attack Known key distinguishing attack	Cryptanalysis technique where the attacker must have access to the communication channel, and obtain the plaintexts corresponding to ciphertexts of their choice. Using this information, the attacker tries to recover the key by breaking an n -bit key cipher into $2^{n/2}$ number of operations.
Timing attack	Cryptanalysis side channel attack where the attacker attempts to break the ciphertext by measuring the execution times of mathematical operations in the encryption process for various inputs.
Birthday attack	Cryptanalysis technique that exploits the mathematics behind the birthday paradox in probability theory, i.e. the higher likelihood of hash collisions found between random attack attempts and a fixed degree of permutations.
Brute force attack	The technique of trying every possible password or key until the correct one is found. May be very time-consuming or even infeasible. Countermeasure: long passwords and keys.
Rainbow table attack	Cracking of password hashes by using a rainbow table, i.e. a pre-computed table that caches the output of cryptographic hash functions. Countermeasure: salt value added to a password before hashing.
Dictionary attack	Password cracking done by trying every entry listed in a dictionary file. The dictionary is a text file containing all words, names, entries with numbers and symbols added to words or replacing letters, finger rolls, etc. Countermeasure: truly random passwords.
Hybrid attack [password cracking]	Password cracking combining a dictionary attack and brute force attack, done by adding numbers and symbols to the dictionary entries.
Meet-in-the-Middle attack	Type of known plaintext attack carried out against ciphers that use multiple keys for encryption. The attacker performs a brute force attack on one key to encrypt the plaintext and on another key to decrypt the ciphertext, trying to find an intermediate ciphertext that matches both.
DUHK (Don't Use Hardcoded Keys)	Cryptographic vulnerability affecting devices that use the ANSI X9.31 Random Number Generator in conjunction with a hardcoded seed key. The attacker can exploit this vulnerability to recover encryption keys from VPN connections or encrypted web sessions.

Virus	Piece of self-replicating code that attaches copies of itself to other executable programs, infecting them.
File virus	Virus that infects an executable file, overwriting it.
Boot Sector virus System virus	Virus that executes its code before the target machine boots. Moves the Master Boot Record to another location on the hard disk and copies itself to the original location of the MBR.
Multipart virus Hybrid virus	Acts both as a file virus and a Boot Sector virus.
FAT virus	Virus that infects the File Allocation Table in FAT filesystems.
Cluster virus	Virus that infects files without changing them, saving instead the virus code to the hard disk and overwriting the pointer in the directory entry; in this way, the virus code is read instead of the actual program.
Stealth virus Tunneling virus	Virus that alters the service call interrupts while running, in order to hide from AV software.
Sparse infector virus	Virus that infects files only occasionally, upon satisfying certain conditions (time, size of the file to infect, etc.) to minimize the risk of being detected by AV software.
Encryption virus	Virus consisting of a decryption module and an encrypted copy of the virus. Once the target machine is infected, the decryption module decrypts and executes the virus. The virus then infects files with a copy of the virus which is encrypted with a different key for each file. This is done to thwart signature detection methods in AV software.
Polymorphic virus	Self-reprogramming virus that changes the malicious code at each infection. Consists of the encrypted virus, a decryption routine, and a mutation engine; the virus code mutates with each infection while the virus algorithm stays the same. Undetectable by AV software.
Metamorphic virus	Self-reprogramming virus that rewrites the code each time it infects a new file: it inserts dead code, reorders instructions, and modifies the program control structure to avoid pattern recognition by AV software. Stealthier than a polymorphic virus.
Cavity virus	Virus that overwrites empty spaces (nulls) in the original file so not to modify its size.
Camouflage virus Companion virus	Virus that creates a companion file with a different file extension for each file to infect; the companion file is executed first and contains the malicious code.
Shell virus	Virus that wraps itself around the infected file, hosting the original program as its subroutine, so that the virus code is executed first.
TSR virus	Terminate and Stay Resident virus. Remains resident in memory after the infected program has terminated execution.
Macro virus	Written as a macro (often in VBA language), infects Microsoft Office files.
File extension virus	Virus that tries to hide itself by adding a fake safe file extension (e.g. TXT) to its executable file.
Logic bomb	Virus that is triggered in response to a specific event.

Trojan horse Trojan	Malicious program packed and concealed, with the help of a wrapper, inside an innocuous program.
Remote Access Trojan (RAT)	Trojan that provides full access to the infected host, including files, shell, screen capture, webcam, microphone, etc.
Backdoor trojan	Trojan that allows bypassing the standard system authentication through IDSs and firewalls.
Botnet trojan	Trojan that infects a large number of hosts (bots aka zombies) to create a botnet, which is then controlled via a Command&Control center to carry out distributed attacks.
Rootkit trojan	Trojan composed of a dropper, a loader, and a rootkit; the dropper runs the loader which causes a buffer overflow, causing the rootkit to be loaded into memory.
Proxy server trojan	Trojan that allows an attacker to use the infected system as a proxy to connect to the Internet.
Covert channel trojan	Trojan that creates a covert channel in the data stream authorized by the network access control system, allowing the attacker to tunnel malicious traffic undetected.

Rootkit	Malicious program used to gain full, administrator-level, persistent access to a system without detection.
Hardware/firmware rootkit	Rootkit located in the firmware (hard disks, BIOS, etc.), creating a persistent malware image.
Bootloader-level rootkit	Rootkit that replaces the bootloader and activates itself before the OS starts.
Kernel-level rootkit	Rootkit running at kernel level in Ring 0 with the highest OS privileges. This is the most difficult type of rootkit to detect.
Hypervisor-level rootkit	Rootkit running in Ring 1, hosting the OS of the target machine as a VM and intercepting all hardware calls made by the target OS.
Library-level rootkit	Rootkit that patches, hooks, or replaces OS system calls with backdoored versions.
Application-level rootkit	Rootkit that replaces application files and modifies process behaviour by injecting malicious code.

Worm	Standalone malicious program that replicates itself and executes through network connections. Usually it does not infect files, and just consumes the CPU and RAM of the infected host.
Keylogger	Program that covertly intercepts and records all keys pressed on the keyboard. May be software or hardware. Most effective when placed in the keyboard hardware or the OS.
Ransomware	Malware that encrypts files in the infected system, preventing the legitimate user from accessing them, and asks the victim for a ransom (usually to be paid online in cryptocurrency).
Spyware	Malware that surreptitiously collects private information and sends it to another entity, compromising the user's privacy and/or the system's security in the process.

A **firewall** monitors incoming and outgoing network traffic, allowing or discarding it based on custom security rules.

Firewall architectures:

Bastion host	Defends the LAN from outside attacks. Has one public interface connected to the Internet, and one private interface connected to the LAN.
Screened subnet	Has one public interface connected to the Internet, one private interface connected to the LAN, and one interface connected to the DMZ. Designed to host servers that offer public services.
Multi-homed firewall	Composed of multiple firewalls, or one firewall with at least three NICs, which allows the protection of multiple network segments.

Firewall technologies:

Packet filtering	Filters packets according to source and destination IP address, source and destination TCP/UDP port, TCP flags, etc.
Circuit-level gateway	Forwards data between networks, making traffic appear to have originated from the gateway.
Application-level filtering	Filters traffic depending on the protocol.
Stateful multilayer inspection	Combines packet filtering, circuit-level gateway, and application-level filtering to remember the state of previous packets and determine whether session packets are legitimate.
Application proxy	Provides protection by preventing direct connections between systems on either side of the firewall, and hence by avoiding exposure of the proxied service. A client connects to the proxy firewall (stateful) which then initiates a new network connection on behalf of the request.
Virtual Private Network (VPN)	Using traffic encryption and encapsulation, provides secure access to a private network through a WAN.
Network Address Translation (NAT)	As in routers, remaps LAN's private IP addresses to Internet-routable public IP addresses.

Techniques for firewall reconnaissance and evasion:

Port scanning	The attacker performs a port scan to identify the firewall type and model.
Firewalking	The attacker sends a TCP/UDP packet with a TTL equal to the number of hops to the firewall plus one, to determine gateway ACL filters and perform network mapping.
IP address spoofing	The attacker spoofs the IP address of a trusted host in the network.
DNS poisoning	The attacker performs DNS poisoning, directing a legitimate user inside the firewalled perimeter to a malicious server which infects the user's host.
Source routing	The attacker sends source routed packets so that they bypass the firewall.
Tiny fragments	The attacker sends tiny fragments of outgoing packets forcing some TCP packet's header information into the next fragment, to bypass user-defined firewall rules.
Proxying	The attacker bypasses URL blacklisting by using an HTTP proxy.
ICMP tunneling	The attacker sends data in the payload portion of an ICMP Echo packet, which is usually not inspected by the firewall.
ACK tunneling	The attacker sends data in a TCP ACK packet, which is usually not inspected by the firewall as it is used as response to legitimate traffic.
HTTP tunneling	The attacker tunnels network traffic through HTTP, to bypass firewalls which only allow HTTP.
SSH tunneling	The attacker tunnels network traffic through SSH, as firewalls often allow SSH.
Encoded XSS	The attacker performs a XSS attack against a WAF (Web Application Firewall), encoding the payload in ASCII or hex to avoid triggering the firewall's filters.

An **IDS (Intrusion Detection System)** monitors, detects, and alerts about possible intrusions (passive mode). An **IPS (Intrusion Prevention System)** also blocks them (active mode). Can be network-based (**NIDS**) or host-based (**HIDS**).

Methods used by an IDS to detect intrusions:

Signature recognition	Pattern-matching of packets over signatures, generated at the network and transport layers, belonging to a known intrusion model.
Anomaly detection	Detection of an event outside the tolerance threshold of normal traffic and behavior.
Protocol anomaly detection	Detection of packets not following protocol standards.

Techniques for IDS evasion:

Insertion attack	The attacker obscures the exploit by inserting extra packets which will be received by the IDS but not by the target host.
TTL attack	Insertion attack in which the attacker sets the appropriate TTL on packets so that they will be received by the IDS but will not reach the target host.
Invalid RST	Insertion attack in which the attacker sends an RST packet with incorrect IP checksum, which the IDS interprets as the end of the session (and hence stops processing the traffic stream), while the target host drops the packet due to invalid checksum.
Urgency flag	Insertion attack in which the attacker sends a packet with the URG flag set. In such a packet the Urgent Pointer field indicates how much of the data in the segment, counting from the first byte, is urgent and should be prioritized. However, unlike the target host, some IDS ignore this field and consider the whole packet payload.
Pre-connection SYN	The attacker sends an initial SYN with invalid TCP checksum before the real connection is established. This desynchronizes the IDS with respect to TCP sequence numbers.
Post-connection SYN	The attacker sends a SYN after the connection is established. This desynchronizes the IDS with respect to TCP sequence numbers, but the target host ignores this packet as it references an already established connection. Then the attacker sends an RST with the correct (for the IDS) TCP sequence number to close its connection.
Evasion	The attacker sends portions of the exploit in crafted packets that the IDS mistakenly does not consider. This can be done e.g. by inserting data in the payload of a TCP SYN.
Unicode evasion	The attacker bypasses the IDS by encoding the packets in Unicode (UTF-8 and UTF-16), which features code points for multiple representations of the same character.
DoS	The attacker overwhelms the IDS capacity (CPU, RAM, disk space, network bandwidth) by sending a large amount of bogus traffic.
Obfuscation	The attacker obscures the exploit by encoding the packet payload in a way that it will be understood by the target host but not by the IDS.
Polymorphic shellcode	The attacker encrypts the packet and adds the code to decrypt it inside the packet, so that the IDS cannot recognize the attack signature. Done via a buffer overflow exploit where the return memory address points to the entrance point of the decryption code. Countermeasure: checking for no-op opcodes other than 0x90.
ASCII shellcode	The attacker encodes the attack code in ASCII characters which translate to assembly instructions, so that the IDS cannot recognize the attack signature.
Application-layer attack	The attacker hides the attack code inside application-layer compressed data (audio, video, images, etc.), so that the IDS cannot recognize the attack signature.
False positive	The attacker sends a large amount of traffic known to trigger alerts on the IDS, hiding the real attack traffic under the noise.
Session splicing	The attacker splits the exploit in several small packets so that no single packet triggers the IDS in itself. The attacker can also add a delay between packets to cause the IDS to timeout before it can reassemble and check them.
Fragmentation attack	If the IDS has a shorter timeout than the target host for reassembling fragmented packets, the attacker sends packet fragments with a long delay; the IDS drops them while the target hosts reassembles them. If the IDS has a longer timeout, the attacker sends mixed valid and bogus packet fragments with a short delay; the IDS gets a reassembled packet with wrong checksum and drops it, while the target host gets the reassembled exploit payload.
Overlapping fragments	The attacker sends the exploit fragmented in packets with overlapping TCP sequence numbers. Different OSes handle fragments differently, so the IDS may end up with a bogus packet while the target host may end up with the exploit payload.

A **WIDS (Wireless Intrusion Detection System)** or **WIPS (Wireless Intrusion Prevention System)** monitors the radio spectrum to detect, block, and alert about possible intrusions in a wireless network.

Threats detected and prevented by a WIDS or WIPS:

Rogue Access Points

Wireless attack tools

DoS attacks

MitM attacks

Honeypots

Evil twin attacks

MAC spoofing

Ad hoc networks

Client misassociation with an AP

Unauthorized association with an AP

Misconfigured APs

A **honeypot** is a part of a computer and network system set up as a bait for attackers. It looks like a legitimate part of the site and appears to contain information or resources valuable to attackers, but in fact it is isolated and heavily monitored. It is used to identify attackers and their techniques, as well as to deflect attacks.

Types of honeypots:

Low-interaction honeypot	Emulates a limited number of services and applications, and provides limited interaction with the attacker. Used to collect information about attack vectors such as network probes and worm infections.
Medium-interaction honeypot	Emulates a real OS with services and applications, but can only respond to preconfigured commands so the attacker might notice that its functioning is limited. Allows gathering more detailed data about an attack.
High-interaction honeypot	Not an emulation, but a real system actually running services and applications, or a whole network architecture of systems (honeynet). As such, it may be compromised or infected by an attacker. All intrusion activities are logged, so it allows collecting a large amount of data about the attacker. A honeywall gateway allows the attacker to interact with the honeynet while limiting outbound traffic via IDS technologies, to control the attacker and preventing harm to computers outside the honeynet.
Production honeypot	Emulates a real production system of an organization. Used to lure attackers to trigger alerts and hence get early warnings about an attack.
Research honeypot	High-interaction honeypot deployed in research, military, or government organizations to get detailed information about cyberattacks in order to study exploits and vulnerabilities.

Characteristics and methods used by honeypots (which can be used to identify them):

Layer 2 tarpit	A tarpit is a service purposely used to delay incoming connections. A Layer 2 tarpit is used to block network penetration of an attacker, but can be detected by the use of the MAC address 00:00:0f:ff:ff:ff which acts as a blackhole for Layer 2 connections.
Layer 4 tarpit	In a Layer 4 tarpit, the honeypot server receives the initial SYN packet from the attacker and replies with a SYN/ACK, but then does not open a socket and forgets everything about the connection. The attacker sends an ACK, believes the TCP three-way handshake to be completed, and starts sending data which never reaches its destination. The attacker does not receive acknowledgment for the sent packets and retransmits them; the connection eventually times out. Layer 4 tarpitting can also be done by setting a TCP Window Size of 0 to block the attacker from sending further data, while keeping the connection open.
Layer 7 tarpit	Layer 7 tarpits are implemented by having services with a very high latency, e.g. an SMTP server which sends useless replies and waits a long time (15 secs or more) between lines.
SYN Proxy	As a protection against SYN flood attacks, a host applying SYN Proxy to a TCP connection responds to the initial SYN packet from the attacker with a specially crafted SYN/ACK, then waits for the ACK in response before forwarding the connection request to the server. The lack of SYN/ACK retransmissions may indicate the presence of a honeypot, especially <code>honeyd</code> .
Bait and switch	Technique consisting in redirecting all network communications to a honeypot once an intrusion attempt is detected. Attacker-side, this can be identified by the change in TCP/IP parameters (RTT, TTL, TCP timestamps, etc.)
Fake Access Points	Honeypots may create fake Wi-Fi APs. These APs have random SSIDs and send beacon frames but no other traffic.
User-mode Linux (UML)	User-mode Linux allows multiple virtual Linux kernel-based guest OSes to run as an application within a host Linux system. It is often used for honeypots, so references to UML in <code>/proc</code> subdirectories may indicate that the attacked system is a honeypot.