

## Processes & Jobs

PROCESSES	
<b>chkconfig</b> [<option> [<service> [<availability>]]]	<p><b>--add</b> &lt;service&gt; adds a service under chkconfig management, <b>--del</b> &lt;service&gt; removes a service, <b>--level</b> &lt;runlevel&gt; specifies the runlevel of the system, <b>--list</b> / <b>--list</b> &lt;service&gt; lists all services / specified service and their/its availability {off   on} on all runlevels</p> <pre># chkconfig httpd on (enables the service at runlevels 2, 3, 4 and 5) # chkconfig --level 234 qemu off (disables the service at runlevels 2, 3, and 4) \$ chkconfig --list   egrep -i "rexec rlogin rsh" (lists the specified services and their availability at all runlevels)</pre>
<b>systemctl</b> <command> [<service ...>] (implemented from RHEL 7)	<p><b>enable</b> enables a service, <b>disable</b> disables a service, <b>is-enabled</b> shows availability of a service ("service unit") automatically after system startup, <b>mask</b> prevents a service from being started automatically or manually, <b>unmask</b> undoes the effect of "mask", <b>--user</b> manages user systemd files (services)</p> <pre># systemctl enable httpd (enables the service automatically after system startup) # systemctl mask postfix.service (prevents the service from being started automatically or manually) \$ systemctl list-unit-files -t service (lists all services and their availability)</pre>
<b>service</b> [<service>] <command>	<p><b>start</b> starts a service, <b>stop</b> stops a service, <b>restart</b> restarts a service, <b>reload</b> reloads a configuration file of the service, <b>status</b> shows the status of the service</p> <pre># service cups restart  # /etc/init.d/cups restart (restarts the specified service) # service --status-all (shows the status of all services)</pre>
<b>systemctl</b> <command> [<service ...>] (implemented from RHEL 7)	<p><b>start</b> starts a service, <b>stop</b> stops a service, <b>restart</b> restarts a service, <b>reload</b> reloads a configuration file of the service, <b>status</b> shows the status of the service ("service unit"), <b>is-active</b> prints whether the service is active (running) or inactive, <b>--user</b> manages user systemd files (services)</p> <pre># systemctl restart cups (restarts the specified service) # systemctl list-units -t service (shows the status of installed services) # systemctl list-units -t service --state=running (shows running services) # systemctl list-units -t service --all (shows the status of all services)</pre>

PROCESSES	
<b>systemctl list-jobs</b> (implemented from RHEL 7)	lists all jobs that systemd is currently executing
<command> [<argument ...>]	executes an executable program (binary file) in the current shell; the program is a shell builtin command or a command found in \$PATH or in the specified path \$ date
<b>bash   sh</b> [<command   file> [<argument ...>]]	executes a shell, command or file (even without execute bits set) in a subshell; the file is searched in the current directory and in \$PATH, <b>-r</b> executes a restricted shell, <b>-c</b> <string> reads commands from the string, <b>-s</b> reads commands from STDIN, <b>-x</b> executes the script in a debug mode \$ sudo bash -c "echo 'Welcome to my website' > /var/www/html/index.html" (redirects the output of the command to the specified file without having to log in to the root user account)
<b>source</b>   . <file> [<argument ...>]	executes a file (even without execute bits set) in the current shell; the file is searched in the current directory and in \$PATH \$ . ~/.bash_profile (refreshes a user profile without the necessity to log off and log on)
<file> [<argument ...>]	executes an executable file in a subshell; the file is searched in \$PATH or in the specified path \$ ./script
<b>command</b> <command> [<argument ...>]	executes a shell builtin command or a command found in \$PATH, ignoring any shell function named the same as the command name (the shell searches for a command name in this order: shell functions, shell builtin commands, executable files in \$PATH)
<b>builtin</b> <command> [<argument ...>]	executes a shell builtin command, ignoring any shell function named the same as the command name (the shell searches for a command name in this order: shell functions, shell builtin commands, executable files in \$PATH)
<b>exec</b> [<command> [<argument ...>]] [<redirection ...>]	replaces the shell with the specified command (no new process is created) or performs redirections \$ exec bash (replaces the shell with bash) \$ exec > output.txt (redirects all output to the "output.txt" file for the current shell process)
<b>ps</b> [<PID   user>]	displays a tree structure of processes in alphabetical order, starting with the "init" (systemd) process or the specified PID or processes owned by the specified user, <b>-n</b> in numerical order (by PID), <b>-p</b> prints PID, <b>-u</b> prints usernames

PROCESSES	
ps	<p>displays a static list of running processes of the logged-in user on the particular terminal (PID - the process identifier, TTY - the terminal, from which the process started (the "?" character means that there is no terminal associated with the process), TIME - the time the process has been processed by the processor, CMD - the command or process and its parameters), <b>-e</b> all processes of all users, <b>-f</b> a detailed output (shows also the UID - the owner of the process, PPID - the parent process identifier, C - the CPU usage, STIME - the time the process started), <b>-l</b> an entire output (shows also F - the process flag (0 - a regular process, 1 - a forked, but not executed process, 4 - a process running with root privileges, 5 - the combination of the previous values), S - the process status ("R" a running process, "S" a sleeping process (waiting for an event to complete), "D" a sleeping process not responding to signals (usually I/O), "T" a stopped (suspended) process, "Z" a zombie (defunct) process - a terminated process that still occupies a PID as the parent process is not aware of its exit status), C - the CPU usage in %, PRI - the priority of the process, NI - the priority value specified by the "nice" command, SZ - the total memory size in blocks occupied by the process, WCHAN - the address of the kernel function where the process waits for a particular event (a running process has "-")), <b>-p</b> &lt;PID&gt; a specified process, <b>-C</b> &lt;program&gt; processes of the specified program, <b>-U</b> &lt;user   UID&gt; processes of a specified user, <b>-G</b> &lt;group   GID&gt; processes of a specified group, <b>-t</b> &lt;terminal&gt; processes associated with a specified terminal, <b>-o</b> &lt;format&gt; according to the specified format, <b>--sort</b> &lt;specification&gt; sorts processes according to the specification, <b>-Z</b> prints the SELinux security context</p> <pre>\$ ps -ef   grep sshd (searches the running processes for the specified process) \$ ps -lt pts/1 (prints a detailed list of running processes on the specified terminal) \$ ps -p \$\$ (prints the current shell process) \$ ps -p 1 -o comm= (prints the name of the process by its PID) \$ ps -eo pid,ppid,user,%mem,%cpu,comm --sort=-%cpu (sorts processes by the CPU usage) \$ ps -eo pid,ppid,user,%mem,%cpu,comm --sort=-%mem (sorts processes by the memory usage)</pre>

PROCESSES	
<b>top</b>	displays a list of processes in an interactive and dynamic way, including their PID, owner, priority or status, and overall CPU, memory and swap usage, <b>-u</b> <user> processes of a specified user, <b>-p</b> <PID> a specified process; interactive options: <b>f</b> displays a menu of items that can be added as an additional column to the command output (e.g., swap, UID or PPID), <b>M</b> sorts processes by memory usage, <b>P</b> sorts processes by processor usage (by default), <b>k</b> <PID> terminates the specified process, <b>q</b> quits the program
<b>pidof</b> <process_name>	prints the PID of a specified process \$ ps -p \$(pidof sshd) (searches the running processes for the specified process)
<b>pgrep</b> <pattern>	prints PIDs of processes matching the pattern, <b>-f</b> matches the full command line (not only the process name), including the executable path, arguments, and options, <b>-l</b> prints the process name as well, <b>-P</b> <PPID> prints processes of the parent process, <b>-c</b> prints the number of matching processes, <b>-i</b> ignores case distinctions, <b>-u</b> <user   UID> processes of a specified user \$ pgrep http (prints PIDs of the httpd processes)
<b>nice</b> [[ <b>-n</b> ] <priority>] <command>	runs a process with modified priority ("nice" value), the "priority" parameter specifies the value to be subtracted from or added to the default value (0 by default), if omitted the value is automatically reduced by 10; the priority is specified in the range between -20 and 19, the lower the number, the higher the priority, negative values can only be set by root; with no argument the current inherited priority is printed \$ nice -n 15 rm -rf ~/tmp/* (decreases the process priority by 15) # nice -n -15 ls /etc   cpio -ov > /dev/rmt0 (increases the process priority by 15)
<b>renice</b> [ <b>-n</b> ] <priority> <identifier>	modifies the priority ("nice" value) of a running process, the value of the priority is specified in the absolute form in the range between -20 and 19; a regular user can only decrease the priority, [ <b>-p</b> ] <PID> specifies a process, <b>-u</b> <user> specifies processes of a particular user, <b>-g</b> <group> specifies processes of a particular group # renice -n 5 987 (sets the priority of the specified process) # renice -n -20 -u root (increases the priority of all processes of the root user to the maximum) # renice -n 19 -g users (decreases the priority of all processes of the "users" group to the minimum)
<b>Ctrl+c</b>	terminates a running process in the foreground
<b>Ctrl+d</b>	terminates a running process that reads data from STDIN or exits the current shell

PROCESSES	
<b>kill</b> [<signal>] <PID ...   %job_number ...>	<p>(<b>-15</b>) terminates a process or job matching the specified ID in a standard way (signal SIGTERM), <b>-9</b> terminates a process immediately (signal SIGKILL), <b>-1</b> terminates a process, for daemons it reloads their configuration file (signal SIGHUP), <b>-17</b> notifies the parent process on child processes events (signal SIGCHLD), <b>-19</b> suspends a process (signal SIGSTOP), <b>-18</b> starts a suspended process (signal SIGCONT), <b>-l</b> lists all signals, <b>-s &lt;signal&gt;</b> specifies a signal; the signal is specified by its number or name (with or without the "SIG" prefix and regardless of case)</p> <pre># kill -9 1 # kill -SIGKILL 1 # kill -KILL 1  # kill -sigkill 1 # kill -kill 1</pre> <p>(terminates a specified process immediately)</p> <pre>\$ kill %3</pre> <p>(terminates a specified job)</p> <pre>\$ kill -s SIGCHLD 1376</pre> <p>(terminates a zombie process by sending the SIGCHLD signal to the parent process)</p>
<b>killall</b> [<signal>] <process_name ...>	<p>(<b>-15</b>) terminates all processes matching the specified name (especially useful for multiple running processes of one program) in a standard way (signal SIGTERM), <b>-9</b> terminates a process immediately (signal SIGKILL), <b>-19</b> suspends a process (signal SIGSTOP), <b>-18</b> starts a suspended process (signal SIGCONT), <b>-l</b> ignores case distinctions, <b>-l</b> lists all signals, <b>-s &lt;signal&gt;</b> specifies a signal, <b>-u &lt;user&gt;</b> terminates all processes of a specified user; the signal is specified by its number or name (with or without the "SIG" prefix and regardless of case)</p> <pre># killall -9 sshd</pre> <p>(terminates specified processes immediately)</p>
<b>pgrep</b> [<signal>] <pattern>	<p>(<b>-15</b>) terminates a process matching the specified pattern in a standard way (signal SIGTERM), <b>-9</b> terminates a process immediately (signal SIGKILL), <b>-19</b> suspends a process (signal SIGSTOP), <b>-18</b> starts a suspended process (signal SIGCONT), <b>-f</b> matches the full command line (not only the process name), including the executable path, arguments, and options, <b>-i</b> ignores case distinctions, <b>-u &lt;user   UID&gt;</b> processes of a specified user; the signal is specified by its number or name (with or without the "SIG" prefix and regardless of case)</p>
<b>fuser</b> [<file ...   device ...>]	<p>prints the processes using a specified file or device, <b>-m</b> prints the processes using the entire file system on which the file resides, <b>-u</b> prints the owner of the process, <b>-k</b> terminates a process, <b>-i</b> asks for confirmation, <b>-v</b> detailed output</p> <pre>\$ fuser -kiv /dev/cdrom</pre> <p>(terminates processes using the specified device)</p>
<b>strace</b> [<command>]	<p>traces system calls and signals called and received by a process, <b>-e &lt;expression&gt;</b> traces a specified system call, <b>-o &lt;file&gt;</b> writes the output to a file, <b>-p &lt;PID&gt;</b> traces a running process, <b>-c</b> counts the time, calls and errors for each system call and reports a summary on program exit</p>

PROCESSES	
<b>ltrace</b> [<command>]	traces shared library calls and signals called and received by a process, <b>-e</b> <expression> traces a specified library call, <b>-o</b> <file> writes the output to a file, <b>-p</b> <PID> traces a running process, <b>-c</b> counts the time and calls for each library call and reports a summary on program exit, <b>-S</b> additionally displays system calls
<b>rpcinfo</b> [<IP_address   hostname>]	<b>-p</b> prints RPC services (their identification number, version, protocol, port and name) running under the portmapper management on the local or specified computer; the list of all RPC programs is stored in <i>/etc/rpc</i>
<b>clustat</b>	displays the status of the cluster
<b>pcs status</b> (implemented from RHEL 6)	displays the status of the cluster

JOBS	
<b>jobs</b> [%<job_number> ...]	prints a specified job (its number, "+" or "-" indicating the most recent or previous job, the status, and the command associated with the job) in the current shell, <b>-l</b> including PID, <b>-r</b> prints only running jobs, <b>-s</b> prints only stopped jobs, <b>-n</b> prints only the changes since the previous execution; with no argument all suspended jobs and jobs running in the background are printed
<command> <b>&amp;</b>	executes a job in the background in a subshell
<b>Ctrl+z</b>	suspends a running job in the foreground
<b>suspend</b>	suspends the execution of the current shell until it receives a SIGCONT signal, <b>-f</b> suspends the login shell
<b>bg</b> [%<job_number> ...]	runs a specified suspended job in the background; with no argument the last suspended job is run
<b>fg</b> [%<job_number>]	runs a specified suspended job in the foreground or moves a specified job from the background to the foreground; with no argument the last suspended job is run or the last running background job is moved to the foreground
<b>nohup</b> <command>	executes a command in such a way that it is resistant to the SIGHUP signal (it runs even after the user has logged out); the output of the command is saved to the "nohup.out" file in the working directory <pre># nohup find /tmp -name core -print &gt; core.txt &amp;</pre> (executes the specified command in the background, which runs even after the user logs out)

JOBS	
<b>watch</b> <command>	<p>executes a command repeatedly at regular intervals of 2 seconds, <b>-n</b> &lt;n&gt; every n seconds, <b>-d</b> highlights the changes in the output, <b>Ctrl+c</b> terminates the program</p> <pre>\$ watch -n 1 date</pre> <p>(displays the current time)</p> <pre>\$ watch -d iostat</pre> <p>(displays changes to disks I/O loads and data write and read speeds)</p> <pre>\$ watch "ps -eo %cpu,pid,user,args --sort=-%cpu   head -11"</pre> <p>(displays a dynamic list of the 10 processes that use the most processor)</p>
<b>wait</b> [<PID   %job_number>]	<p>waits for a specified process or job running in the background of the current shell to finish, then terminates; it returns the last process termination status, if the process doesn't exist, the return code is 127; with no argument it waits for all the current shell background processes to terminate and the return code is 0</p> <pre>\$ wait \$! &amp;&amp; echo "OK"</pre> <p>(waits for the last process running in the background to complete)</p>
<b>sleep</b> <duration>	<p>sets a time interval during which no activity is performed (especially useful when a certain delay is needed before executing the next command); the duration is specified by a number and a suffix <b>s</b> in seconds (by default), <b>m</b> in minutes, <b>h</b> in hours, <b>d</b> in days</p> <pre>\$ sleep 10 &amp;&amp; echo "10 seconds have passed."</pre> <p>(executes the following command after 10 seconds)</p>
<b>at</b> <time> [<date>] [<additional_time_specification>]	<p>executes a one-time job at the specified time; the commands to be run are read from STDIN or from a file, <b>-f</b> &lt;file&gt; reads commands from a specified file, <b>-c</b> &lt;job_number&gt; displays the contents of the scheduled job, <b>-d</b> &lt;job_number&gt; removes a scheduled job (equivalent to the "atrm" command), <b>-l</b> prints a list of scheduled jobs (equivalent to the "atq" command), <b>-m</b> sends an email to the user when the job has completed</p> <pre>\$ at 20:00 25.06.2007 &lt;-'   \$ at 7am today + 3 weeks &lt;-'</pre> <pre>at&gt; mail root &lt; offer.txt &lt;-'</pre> <pre>at&gt; Ctrl+d</pre> <p>(sends an email to the root user at the particular time with the contents of the specified file)</p> <pre>\$ echo "date &gt; date.txt"   at noon Sunday</pre> <p>(executes a command at the particular time, the output of which is written to the specified file)</p> <pre>\$ at -mf list 20:00 25.06.2007</pre> <p>(executes a list of commands specified in the file at the particular time and sends an email to the user when the job has completed)</p>
<b>atq</b>   <b>at -l</b>	prints a list of scheduled jobs (a job number, time of execution, a user who scheduled the job)
<b>atrm</b>   <b>at -d</b> <job_number ...>	removes a scheduled job

JOBS	
<b>lpr</b> [<file>]	prints a file; if no file is specified, reads from STDIN, <b>-P</b> <printer> specifies a printer, <b>-#</b> <n> specifies the number of copies
<b>lpq</b>	displays jobs in the print queue of the default printer, <b>-P</b> <printer> specifies a printer, <b>-a</b> jobs on all printers, <b>-l</b> detailed output
<b>lprm</b> [<job_number>   -]	removes a specified job from the print queue or the entire queue, <b>-P</b> <printer> specifies a printer; with no argument the current job on the default printer is removed
<b>crontab</b> [<file>]	<p>configures an existing file to periodically execute the specified jobs via the cron daemon, <b>-e</b> creates or edits a crontab file, <b>-l</b> prints the contents of the crontab file, <b>-r</b> removes a regular crontab file, <b>-u</b> &lt;user&gt; sets up a crontab file of the specified user; the jobs are specified each on a separate line in the file, which in the case of user files consists of 6 fields separated by spaces - 5 for time specification (in the order of minute, hour, day of the month, month, and day of the week (Sun-Sat = 0-6), "*" matches all values, "*" &lt;n&gt; every nth time interval, "," separates values of the same field, "-" indicates a range of values) and 1 for the command itself (the command must be specified using the absolute path, except for shell builtin commands and keywords, if it has any output, it is sent to the user by email); an empty line and a line starting with a "#" character are ignored; system files include <i>/etc/crontab</i> and files located in <i>/etc/cron.d</i>, which contain 1 extra field between the time specification and the command - the user under which the specified job should be run; user files are located in <i>/var/spool/cron/&lt;user&gt;</i> and if they are edited manually (not with the "crontab -e" command), the changes are not automatically loaded by the daemon; authorized users are specified in <i>/etc/cron.allow</i>, unauthorized users in <i>/etc/cron.deny</i>, if both the files exist, <i>/etc/cron.deny</i> is ignored</p> <pre>30 8 * * * /usr/bin/df   /bin/mail admin (executes a job every day at 8:30) */30 8-16 * * 1-5 /usr/bin/who &gt;&gt; /tmp/users.out (executes a job every 30 minutes in working hours) 0 22 * * 5 /usr/local/scripts/backup.sh (executes a script every Friday at 22:00) 0 0 1 * * for file in \$(/bin/find /web/logs*); do &gt; \$file; done (executes a job at the beginning of each month)</pre>

From: <https://prompt.cz/en/> - **Prompt.cz**

Permanent link: <https://prompt.cz/en/processes-and-jobs>

Last update: **2025/07/18 12:48**

