

Introduction to Linux

Giorgio Agugiaro

Last update: 22 October 2024

License

This presentation is licensed under the [Creative Commons License CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/). According to CC BY-NC-SA 4.0 permission is granted to share this document, i.e. copy and redistribute the material in any medium or format, and to adapt it, i.e. remix, transform, and build upon the material under the following conditions:



- **Attribution:** You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.
- **NonCommercial:** You may not use the material for commercial purposes.
- **ShareAlike:** If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.
- **No additional restrictions:** You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Index

- Linux in a nutshell
- The terminal and the shell(s)
- File system
- File system permissions
- Data streams
- Processes
- Scheduling jobs
- Installation of software applications
- Text editors
- Bash shell scripting
- Further resources

Index

- **Linux in a nutshell**
- The terminal and the shell(s)
- File system
- File system permissions
- Data streams
- Processes
- Scheduling jobs
- Installation of software applications
- Text editors
- Bash shell scripting
- Further resources

Linux in a nutshell

Linux in a nutshell

Terminal & shell(s)
File system
File system
permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
Bash shell scripting
Further resources

- **Linux** (more precisely: **GNU/Linux**¹) is a computer operating system that uses the GNU software *and* the Linux kernel
- **GNU** is a collection of free and open software packages which can be *either* used as a stand-alone operating system *or* can be used in parts in other operating systems – as in the case of GNU/Linux. [Richard Stallman](#) started the development of GNU in 1984
 - Examples: GCC (GNU C Compiler), GNU Bash shell, etc.
- The [Linux kernel](#) was originally developed by [Linus Torvalds](#) and first released in 1991. A kernel is a computer program, always loaded in memory, and is the core of a computer's operating system. It controls all resources and applications, and it facilitates the interaction between hardware and software.



¹ If you are curious, you can read here more on the long-standing naming controversy:
https://en.wikipedia.org/wiki/GNU/Linux_naming_controversy

Linux in a nutshell

Linux is used in personal computers, but not only!

You'll find it, among the rest, in:

- all Android devices
- routers, NAS, TVs, eReaders, smart watches, cars...
- the majority of servers, world-wide (96%)
- high-performance computers, world-wide (nearly 100%)
- ...the International Space Station!! 😊

Linux in a nutshell

Terminal & shell(s)

File system

File system

permissions

Data streams

Processes

Scheduling jobs

Software install

Text editors

Bash shell scripting

Further resources



Image source: <https://www.asi.it/wp-content/uploads/2019/03/iss.jpg>

Linux in a nutshell

Linux in a nutshell

Terminal & shell(s)

File system

File system

permissions

Data streams

Processes

Scheduling jobs

Software install

Text editors

Bash shell scripting

Further resources

- Linux comes packaged in a so-called **Linux distribution** (short: "distro") that assembles the kernel, the GNU tools and a selection of several other software packages tailored for specific needs
- Some common distributions are Debian, Fedora, openSUSE, ArchLinux, Gentoo. Other distros "derive" from the previous ones: Ubuntu, Red Hat Enterprise Linux, SUSE Linux Enterprise, Manjaro Linux, etc.



- Some distros are tailored to servers, other to personal computers. The latter ones generally come with a Desktop Environment (DE), consisting of several programs with a Graphical User Interface (GUI)
- There are different **Desktop Environments**. Among the most common ones are KDE Plasma, GNOME, MATE, LXDE, etc.



Just like in other operating systems, there are **two types of users**. Each type of users is granted different privileges. They are:

- **root**: is the "name" of the main *superuser* in Linux. A superuser is a special account user used for system administration. The "root" user has all rights and permissions to all files and programs in all modes (single- or multi-user). The "root" user should never ever be used to perform operations on the computer other than system administration!
- **"Normal" users**: are all other users that are not "root". They have limited privileges on what files/directories they can read and write, and what programs they can run. Normal users are associated to a **user name**, such as "giorgio", "luke", "ashoka", etc, and to a **user group**.
 - In many Linux distributions (like Ubuntu), the **sudo** command enables "normal" users to run programs with the security privileges of another user, by default the superuser "root".
For example: updating the software requires the following two commands to be run by the "root" user – unless a "normal" user adds the sudo command as follows:

sudo apt update

sudo apt upgrade

Working with Linux

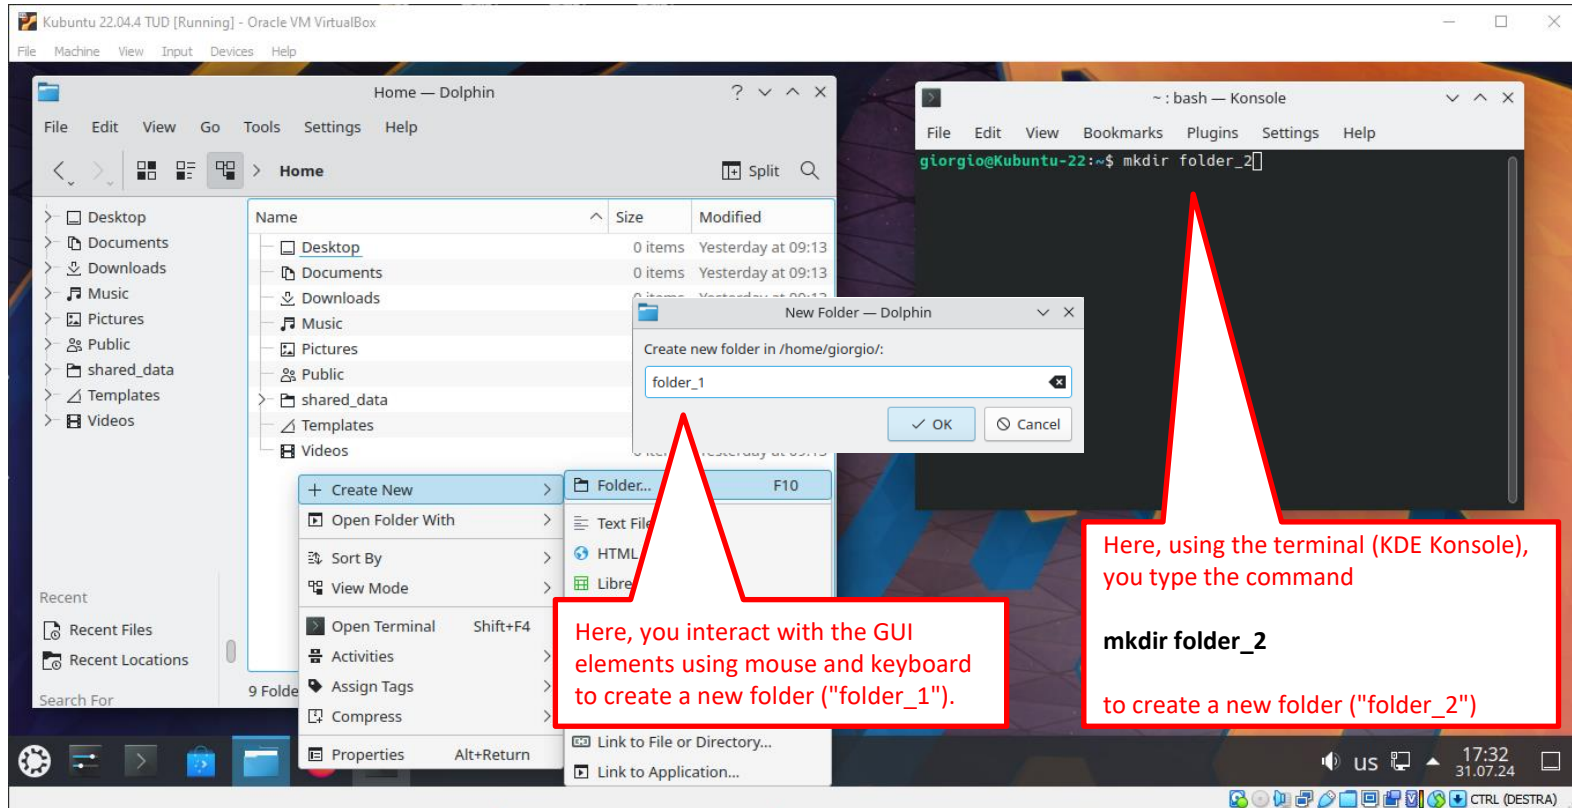
- Just like in other operating systems, you can interact with Linux in two ways:
 - Using the **terminal** to write commands that the system will carry out
 - Generally faster, but requires to know the command names and their syntax
 - You can cluster and run multiple commands together by writing and running a script
 - Using the **GUI elements** (windows, buttons, etc.)
 - More user-friendly, commands can be run clicking buttons or interacting with the GUI
 - It may be more difficult to automatize repetitive operations using the GUI elements (some commands may not have a GUI at all...)
 - GUI elements may not always be available (e.g. working remotely on a server)
- See next slide for an example

Linux in a nutshell
Terminal & shell(s)
File system
File system
permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
Bash shell scripting
Further resources

Working with Linux

Example: Create a new folder in your home directory

Linux in a nutshell
Terminal & shell(s)
File system
File system
permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
Bash shell scripting
Further resources



The screenshot shows a KDE desktop environment with two windows. The Dolphin file manager window is open to the Home directory, displaying a list of folders. A 'New Folder' dialog box is open, showing the folder name 'folder_1' in the input field. A context menu is also visible over the Dolphin window, with 'Folder...' selected. The Konsole terminal window is open, showing the command `mkdir folder_2` being entered. Red callout boxes highlight the GUI interaction and the terminal command.

Here, you interact with the GUI elements using mouse and keyboard to create a new folder ("folder_1").

Here, using the terminal (KDE Konsole), you type the command **mkdir folder_2** to create a new folder ("folder_2")

Index

- Linux in a nutshell
- **The terminal and the shell(s)**
- File system
- File system permissions
- Data streams
- Processes
- Scheduling jobs
- Installation of software applications
- Text editors
- Bash shell scripting
- Further resources

The terminal and the shell(s)

Linux in a nutshell

Terminal & shell(s)

File system

File system

permissions

Data streams

Processes





Scheduling jobs

Software install

Text editors

Bash shell scripting

Further resources

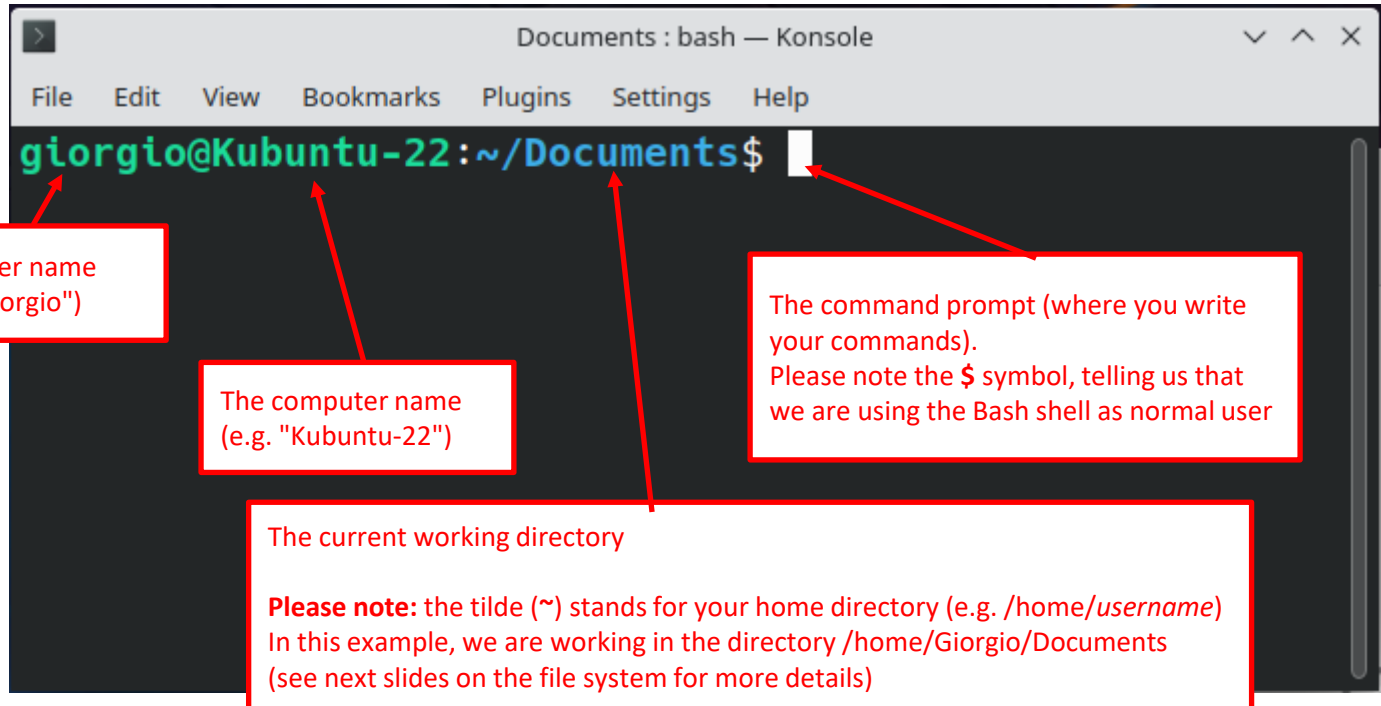
- The **terminal** (also called the **console**, or **command line interface (CLI)**) is a tool to interact with computers by typing textual commands on your keyboard. A terminal is the "window" in which you type commands. It handles user input and output
- A terminal uses a **shell**. A shell is a program that takes the commands you type and translates them into actions that the operating system has to perform. There are several shells, the most common ones are:
 - **bash** (Bourne-Again Shell): the most common one on Linux (and used in this guide) 
 - the bash prompt is \$ for a regular user and # for root (see next slide)
 - **zsh** (z shell): Extended bash with many improvements 
 - the zsh prompt is % for a regular user and # for root
 - **csh** (C-shell): It mimic the C language as the Linux kernel is predominantly written in C 
 - the csh prompt is % for a regular user and # for root
 - **ksh** (KornShell): implements and extends features from the C shell and Bourne shell 
 - the ksh prompt is \$ for a regular user and # for root
 - ...more shells can be found [here](#)

Please note: while the terms "terminal" and "shell" are often colloquially used interchangeably, they are not the same!

The terminal and the Bash shell

Whenever you open a terminal, you get the following information when using the Bash shell:

Linux in a nutshell
Terminal & shell(s)
File system
File system permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
Bash shell scripting
Further resources



Documents : bash — Konsole

File Edit View Bookmarks Plugins Settings Help

giorgio@Kubuntu-22:~/Documents\$

Your user name (e.g. "giorgio")

The computer name (e.g. "Kubuntu-22")

The current working directory

The command prompt (where you write your commands). Please note the \$ symbol, telling us that we are using the Bash shell as normal user

Please note: the tilde (~) stands for your home directory (e.g. /home/username) In this example, we are working in the directory /home/Giorgio/Documents (see next slides on the file system for more details)

Shell commands

Most of the shell commands have a common structure, i.e:

command + options + arguments

Example: The next command copies the `source_directory` and all its contents to the `dest_directory`.

```
cp -rf ./source_directory ./dest_directory
```

- **command:** cp (copy)
- **options:** -rf = -r recursively, -f force
- **arguments:** ./source_directory ./dest_directory

If you do not remember the correct syntax? No problem!

- Most of the commands have a "--help" option
 - Examples: `cp --help`, `rm --help`, `mkdir --help`
- The `man` ("manual") command loads and shows the extensive manual
 - Example: `man cp` will load the on-line manual of the "cp" command

Linux in a nutshell

Terminal & shell(s)

File system

File system

permissions

Data streams

Processes

Scheduling jobs

Software install

Text editors

Bash shell scripting

Further resources

Some useful shell commands

These are very common Bash commands

- `exit`: exit and close the terminal window
- `echo <text>`: displays a line of text
- `cat <file>` ("concatenate"): print file (or group of files) to screen
- `head <file>`: print the first lines of a file
- `tail <file>`: print the last lines of a file
- `which <command>`: locate a command
- `whereis <command>`: locate the binary, source, and manual page files for a command
- `locate <file>`: find file(s) by name, quickly, generally using an index created/updated by `updatedb`
- `touch <file>`: create a new empty file

With many commands you can use also **wildcards**, such as:

- `?`: Matches any single character
- `*`: Matches any string of characters
- `[set]`: Matches any character in the set. Example: `[adf]` will match any occurrence of a, d, f
- `[!set]`: Matches any character NOT in the set of characters

Examples

- `ls *.txt` -> list all files having the extension like "txt"

Linux in a nutshell

Terminal & shell(s)

File system

File system

permissions

Data streams

Processes

Scheduling jobs

Software install

Text editors

Bash shell scripting

Further resources

Index

- Linux in a nutshell
- The terminal and the shell(s)
- **The file system**
- File system permissions
- Data streams
- Processes
- Scheduling jobs
- Installation of software applications
- Text editors
- Bash shell scripting
- Further resources

File system

- In Linux (and other Unix-like operating systems) files are written in a **hierarchical file system**
- Such file system is standardised and is called **Filesystem Hierarchy Standard (FHS)**
- In the FHS, all files and directories appear under the **root directory /**, even if they are stored on different physical or virtual devices
- The first-level directories names (**/bin**, **/usr**, **/home**, **/var**) are consistent over the Unix-like operating systems, and are generally used in the same way

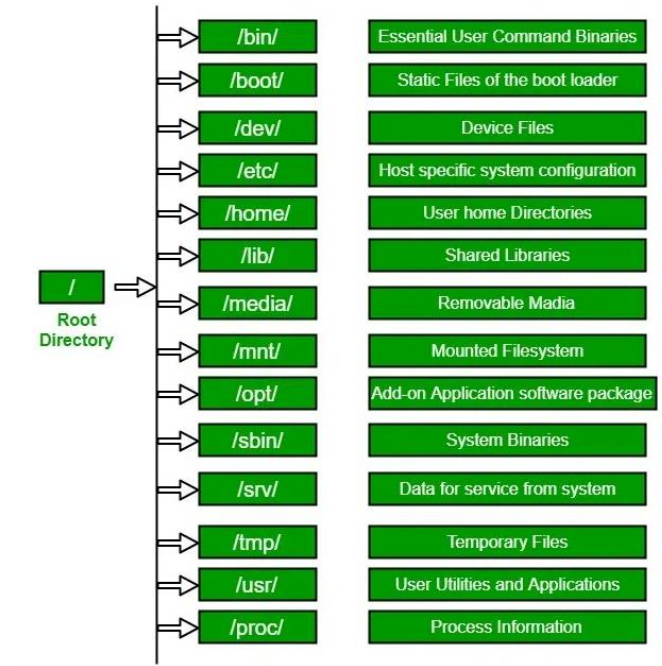


Image source: <https://www.geeksforgeeks.org/linux-file-hierarchy-structure>

Linux in a nutshell
Terminal & shell(s)
File system
File system
permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
Bash shell scripting
Further resources

File system

Some of these directories are particularly relevant:

- **/ (aka "root directory")**: Base directory of the entire file system hierarchy
- **/root**: the "root" user's home directory
- **/home**: all users have their own home directory. Normal users can write ONLY inside their home directory (and subdirectories).
 - Example: /home/giorgio, /home/leia, /home/kylo, etc.
- **/sbin** and **/usr/sbin**: contains program executables ("binaries") that are used for system administration and can be run only by a user with superuser privileges
- **/bin** and **/usr/bin**: contains the majority of the binaries that are installed by default and that can be run by "normal" users.
- **/usr** (and subdirectories): contains the majority of the user utilities and applications
- **/lib**: software libraries necessary for the binaries in /bin and /usr/bin
- **/media**: mount points for removable devices (USB sticks, CD-ROMs, etc.)
- **/mnt**: temporarily mounted file systems
- **/dev/null**: is a special device, used to dispose unwanted output streams of a process, or as a convenient empty file for input streams. This is usually done by *redirection* (see later for examples)

Linux in a nutshell
Terminal & shell(s)

File system

File system

permissions

Data streams

Processes

Scheduling jobs

Software install

Text editors

Bash shell scripting

Further resources

File system

Some common Bash commands to work with the file system

- `pwd` ("print working directory"): tell you in which directory in which you currently are
- `tree`: show all subdirectories from the directory you are in
- `ls` ("list short"): list files and subdirectories in the directory you are in (no details)
- `ll` ("list long", alternative for `ls -l`): list files and subdirectories in the directory you are in (with all details)
- `cd <directory>` ("change directory"): change to another directory
- `mkdir <directory>` ("make directory"): create a new directory
- `rmdir <directory>` ("remove directory"): remove an (empty) directory
- `cp <source-file> <destination-file>` ("copy"): create a copy of a file
- `rm <file> or <directory>` ("remove"): remove a file or a populated directory
- `mv <file> <new-destination>` ("move"): move a file to another directory

Navigation in the file system

- `.` (dot): is the current directory
- `..` (two dots): is the parent directory
- `~` (tilde): stands for your own home directory

Linux in a nutshell
Terminal & shell(s)

File system

File system

permissions

Data streams

Processes

Scheduling jobs

Software install

Text editors

Bash shell scripting

Further resources

File system

Examples:

```
Documents : bash — Konsole
File Edit View Bookmarks Plugins Settings Help
giorgio@Kubuntu-22:~/Documents$ pwd
/home/giorgio/Documents
giorgio@Kubuntu-22:~/Documents$ touch test_file.txt
giorgio@Kubuntu-22:~/Documents$ ls
test_file.txt
giorgio@Kubuntu-22:~/Documents$ cd ..
giorgio@Kubuntu-22:~$ ls
Desktop  Downloads  Pictures  snap      Videos
Documents Music      Public   Templates
giorgio@Kubuntu-22:~$ mkdir test_directory
giorgio@Kubuntu-22:~$ ls
Desktop  Downloads  Pictures  snap      test_directory
Documents Music      Public   Templates Videos
giorgio@Kubuntu-22:~$ rmdir test_directory
giorgio@Kubuntu-22:~$ ls
Desktop  Downloads  Pictures  snap      Videos
Documents Music      Public   Templates
giorgio@Kubuntu-22:~$ cd ./Documents/
giorgio@Kubuntu-22:~/Documents$ pwd
/home/giorgio/Documents
giorgio@Kubuntu-22:~/Documents$
```

```
/: bash — Konsole
File Edit View Bookmarks Plugins Settings Help
giorgio@Kubuntu-22:/$ tree -L 1
.
├── bin -> usr/bin
├── boot
├── cdrom
├── dev
├── etc
├── home
├── lib -> usr/lib
├── lib32 -> usr/lib32
├── lib64 -> usr/lib64
├── libx32 -> usr/libx32
├── lost+found
├── media
├── mnt
├── opt
├── proc
├── root
├── run
├── sbin -> usr/sbin
├── snap
├── srv
├── swapfile
├── sys
├── tmp
├── usr
└── var
```

From root directory /,
print all subdirectories
of first level

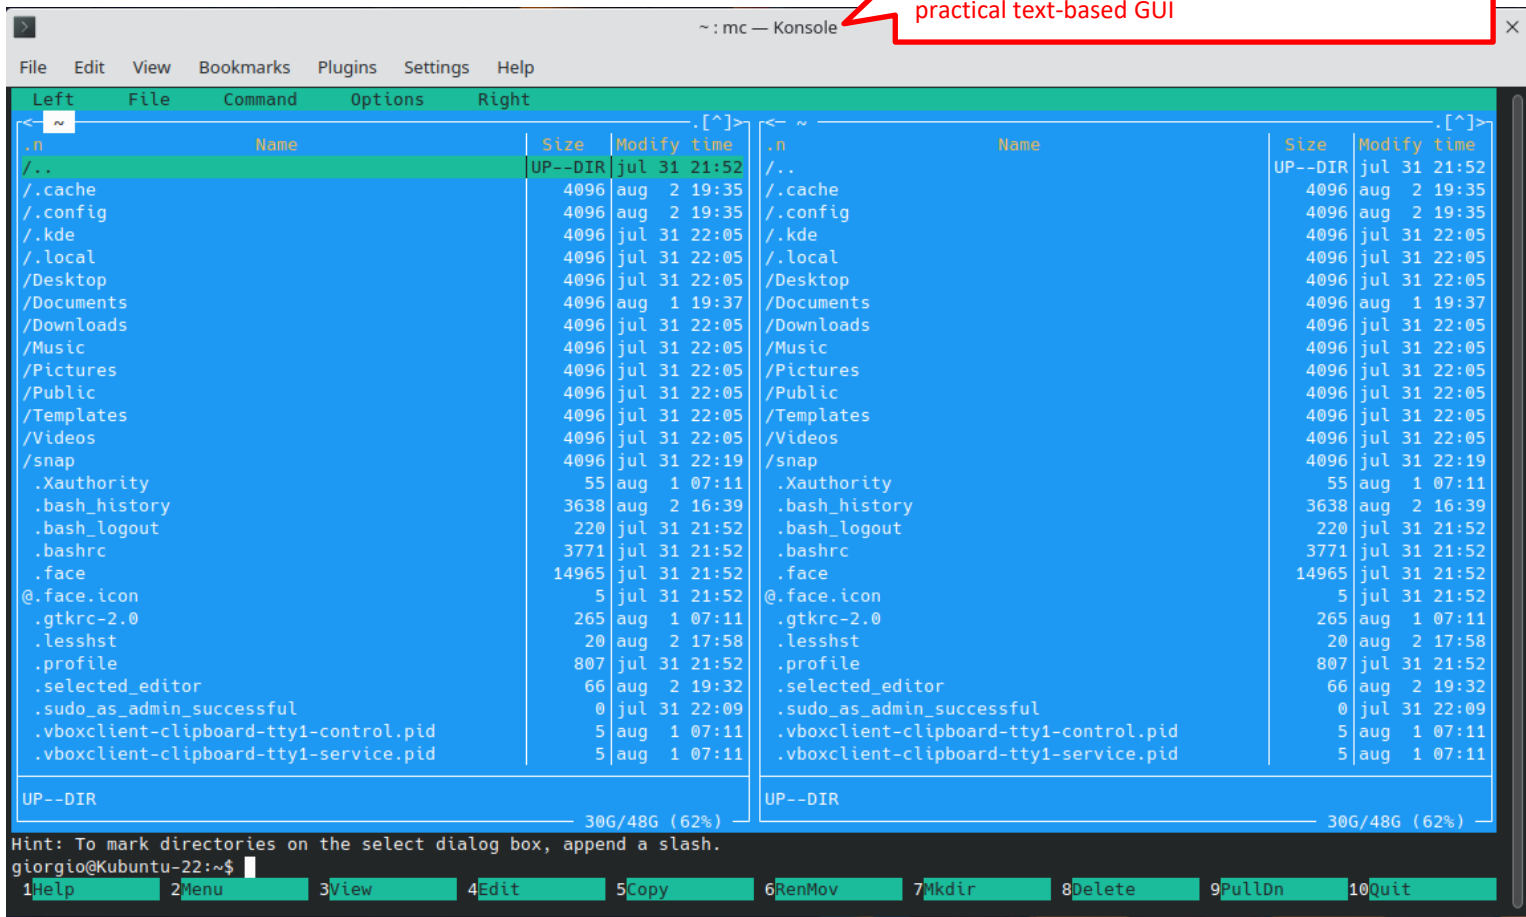
Perform a series of operations using **pwd**, **touch**, **cd**, **mkdir**, **ls**, **rmdir** commands.
You can see the output right after the command has been issued

- Linux in a nutshell
- Terminal & shell(s)
- File system**
- File system
- permissions
- Data streams
- Processes
- Scheduling jobs
- Software install
- Text editors
- Bash shell scripting
- Further resources

File system

Alternatively, you can use for example **mc** ("**M**idnight **C**ommander"), which is a very powerful file manager with comes with a simple but very practical text-based GUI

- Linux in a nutshell
- Terminal & shell(s)
- File system**
- File system
- permissions
- Data streams
- Processes
- Scheduling jobs
- Software install
- Text editors
- Bash shell scripting
- Further resources



The screenshot shows the Midnight Commander (mc) interface in a terminal window. The window title is "~ : mc — Konsole". The menu bar includes File, Edit, View, Bookmarks, Plugins, Settings, and Help. The interface is split into two panes, Left and Right, each displaying a directory listing. The top bar shows "Left File Command Options Right".

| Left | File | Command | Options | Right |
|--|--|---------|-------------|--|
| .n | Name | Size | Modify time | .n |
| UP--DIR | UP--DIR | UP--DIR | UP--DIR | UP--DIR |
| .. | .. | 4096 | aug 2 19:35 | .. |
| /.cache | /.cache | 4096 | aug 2 19:35 | /.cache |
| /.config | /.config | 4096 | aug 2 19:35 | /.config |
| /.kde | /.kde | 4096 | aug 2 19:35 | /.kde |
| /.local | /.local | 4096 | aug 2 19:35 | /.local |
| /Desktop | /Desktop | 4096 | aug 2 19:35 | /Desktop |
| /Documents | /Documents | 4096 | aug 2 19:35 | /Documents |
| /Downloads | /Downloads | 4096 | aug 2 19:35 | /Downloads |
| /Music | /Music | 4096 | aug 2 19:35 | /Music |
| /Pictures | /Pictures | 4096 | aug 2 19:35 | /Pictures |
| /Public | /Public | 4096 | aug 2 19:35 | /Public |
| /Templates | /Templates | 4096 | aug 2 19:35 | /Templates |
| /Videos | /Videos | 4096 | aug 2 19:35 | /Videos |
| /snap | /snap | 4096 | aug 2 19:35 | /snap |
| .Xauthority | .Xauthority | 55 | aug 1 07:11 | .Xauthority |
| .bash_history | .bash_history | 3638 | aug 2 16:39 | .bash_history |
| .bash_logout | .bash_logout | 220 | aug 2 16:39 | .bash_logout |
| .bashrc | .bashrc | 3771 | aug 2 16:39 | .bashrc |
| .face | .face | 14965 | aug 2 16:39 | .face |
| @.face.icon | @.face.icon | 5 | aug 2 16:39 | @.face.icon |
| .gtkrc-2.0 | .gtkrc-2.0 | 265 | aug 2 16:39 | .gtkrc-2.0 |
| .lessht | .lessht | 20 | aug 2 16:39 | .lessht |
| .profile | .profile | 807 | aug 2 16:39 | .profile |
| .selected_editor | .selected_editor | 66 | aug 2 16:39 | .selected_editor |
| .sudo_as_admin_successful | .sudo_as_admin_successful | 0 | aug 2 16:39 | .sudo_as_admin_successful |
| .vboxclient-clipboard-tty1-control.pid | .vboxclient-clipboard-tty1-control.pid | 5 | aug 2 16:39 | .vboxclient-clipboard-tty1-control.pid |
| .vboxclient-clipboard-tty1-service.pid | .vboxclient-clipboard-tty1-service.pid | 5 | aug 2 16:39 | .vboxclient-clipboard-tty1-service.pid |
| UP--DIR | UP--DIR | UP--DIR | UP--DIR | UP--DIR |

At the bottom of the terminal, there is a prompt: `giorgio@Kubuntu-22:~$`. Below the terminal window, there is a navigation bar with the following items: 1Help, 2Menu, 3View, 4Edit, 5Copy, 6RenMov, 7Mkdir, 8Delete, 9PullDn, 10Quit.

Index

- Linux in a nutshell
- The terminal and the shell(s)
- File system
- **File system permissions**
- Data streams
- Processes
- Scheduling jobs
- Installation of software applications
- Text editors
- Bash shell scripting
- Further resources

File system permissions

Every file is described by a 10-character string (called **mode string**) in which:

- Character 1: type of file
- Characters 2-4: privileges of the file owner on that file
- Characters 5-7: privileges of the owner's group on that file
- Characters 8-10: privileges of "everybody else" on that file

Type of file can be:

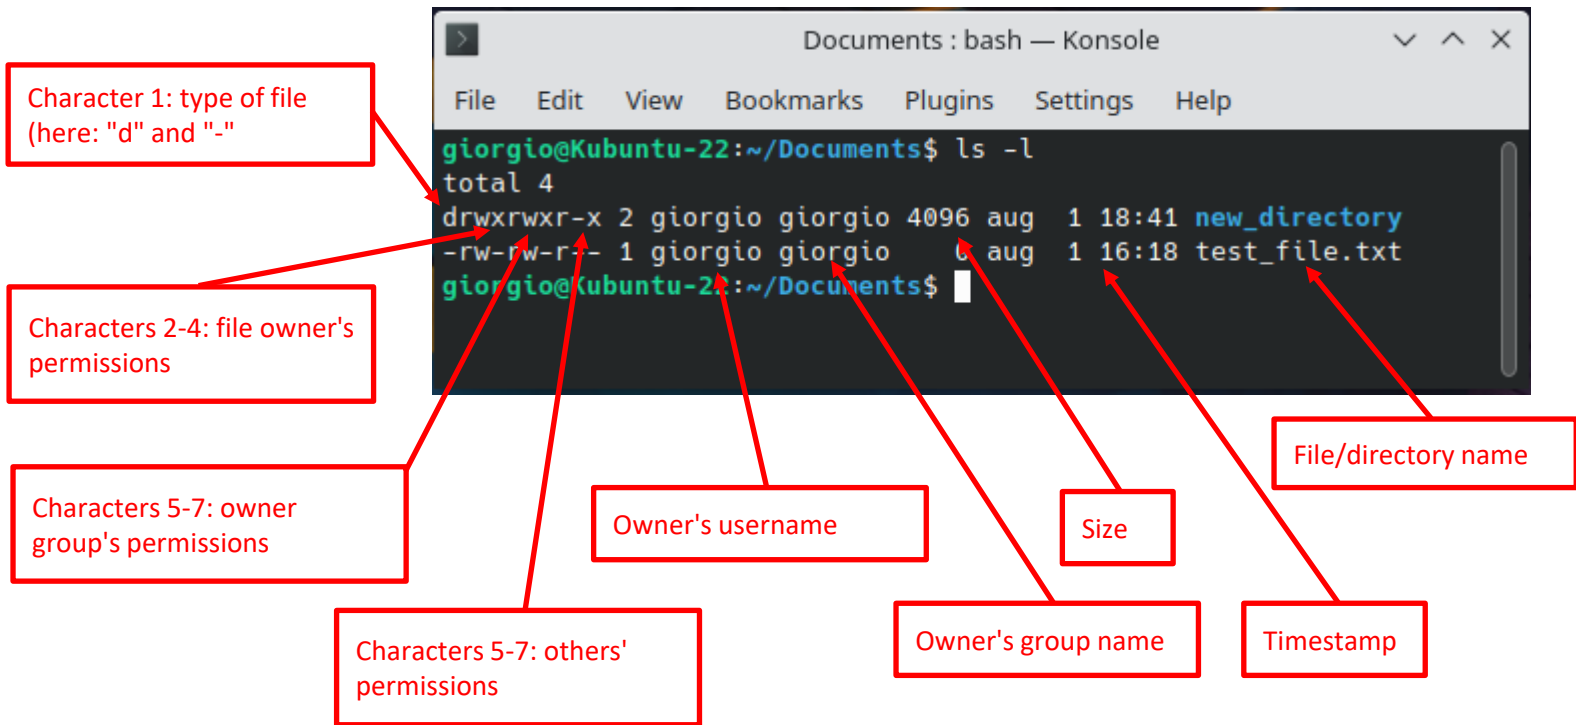
- -: a file
- d: a directory
- l: a symbolic link
- b: a block special file or block device (e.g. /dev/hda, a hard disk)
- c: character special file (e.g. /dev/tty, the terminal of the current process)
- p: a pipe (a temporary file between two linked commands – see later for more examples)
- s: a socket

Type of permission can be:

- r: readable, -: it is not readable
- w: writable, -: it is not writable
- x: executable, or permission to enter a directory, -: it is not executable

File system permissions

Using the command `ls -la` (or `ll -a`), you get all details about each file/directory in the current directory (including the hidden files, with the `-a` parameter).



The terminal window shows the output of the `ls -l` command in a directory named `~/Documents`. The output lists two items: a directory named `new_directory` and a file named `test_file.txt`. Red callout boxes with arrows point to specific parts of the output to explain their meaning:

- Character 1:** type of file (here: "d" and "-")
- Characters 2-4:** file owner's permissions
- Characters 5-7:** owner group's permissions
- Characters 5-7:** others' permissions
- Owner's username:** giorgio
- Owner's group name:** giorgio
- Size:** 4096 (for the directory) and 6 (for the file)
- Timestamp:** aug 1 18:41 (for the directory) and aug 1 16:18 (for the file)
- File/directory name:** new_directory and test_file.txt

File system permissions

The types of permissions can be also expressed numerically. In this way, all possible combinations can be expressed with a digit between 0 and 7.

- r: read = 4
- w: write = 2
- x: execute = 1

Both representations, literal and numeric, are commonly used, especially with **chmod** ("change mode"), a command used to change the permissions of a file or a directory.

| # | Sum | rwX | Permission |
|---|--------------------|-----|-------------------------|
| 7 | 4(r) + 2(w) + 1(x) | rwx | read, write and execute |
| 6 | 4(r) + 2(w) | rw- | read and write |
| 5 | 4(r) + 1(x) | r-x | read and execute |
| 4 | 4(r) | r-- | read only |
| 3 | 2(w) + 1(x) | -wx | write and execute |
| 2 | 2(w) | -w- | write only |
| 1 | 1(x) | --x | execute only |
| 0 | 0 | --- | none |

File system permissions

The command [chmod](#) can be issued in several ways. The permissions to add, remove or change can be expressed using either numerical or symbolic modes.

Here are some examples for using the numerical mode.

- `chmod 664 file_name.txt`: file_name.txt will receive read and write (6) permissions for both the owner and the owner's group, and only read permissions for the "others"
- `chmod 700 file_name.txt`: file_name.txt will receive read, write and execution (7) permissions for the owner and zero permission for the owner's group and the "others"

Please refer to the manual, or **--help** for more details

File system permissions

The command **chmod** can be issued in several ways. The permissions to add, remove or change can be expressed using either numerical or symbolic modes.

The symbolic mode is composed of three components, which are combined to form a single string of text. Specific modes can be modified, leaving the others untouched

```
chmod [references][operator][modes] file
```

| Reference | Class | Description |
|-----------|---------|---|
| u | user | file owner |
| g | group | members of the file's group |
| o | others | users who are neither the file's owner nor members of the file's group |
| a | all | all three of the above, same as ugo |
| (empty) | default | same as "all", except that bits in the <code>umask</code> will be unchanged |

| Operator | Description |
|----------|--|
| + | adds the specified modes to the specified classes |
| - | removes the specified modes from the specified classes |
| = | the modes specified are to be made the exact modes for the specified classes |

Source: <https://en.wikipedia.org/wiki/Chmod>

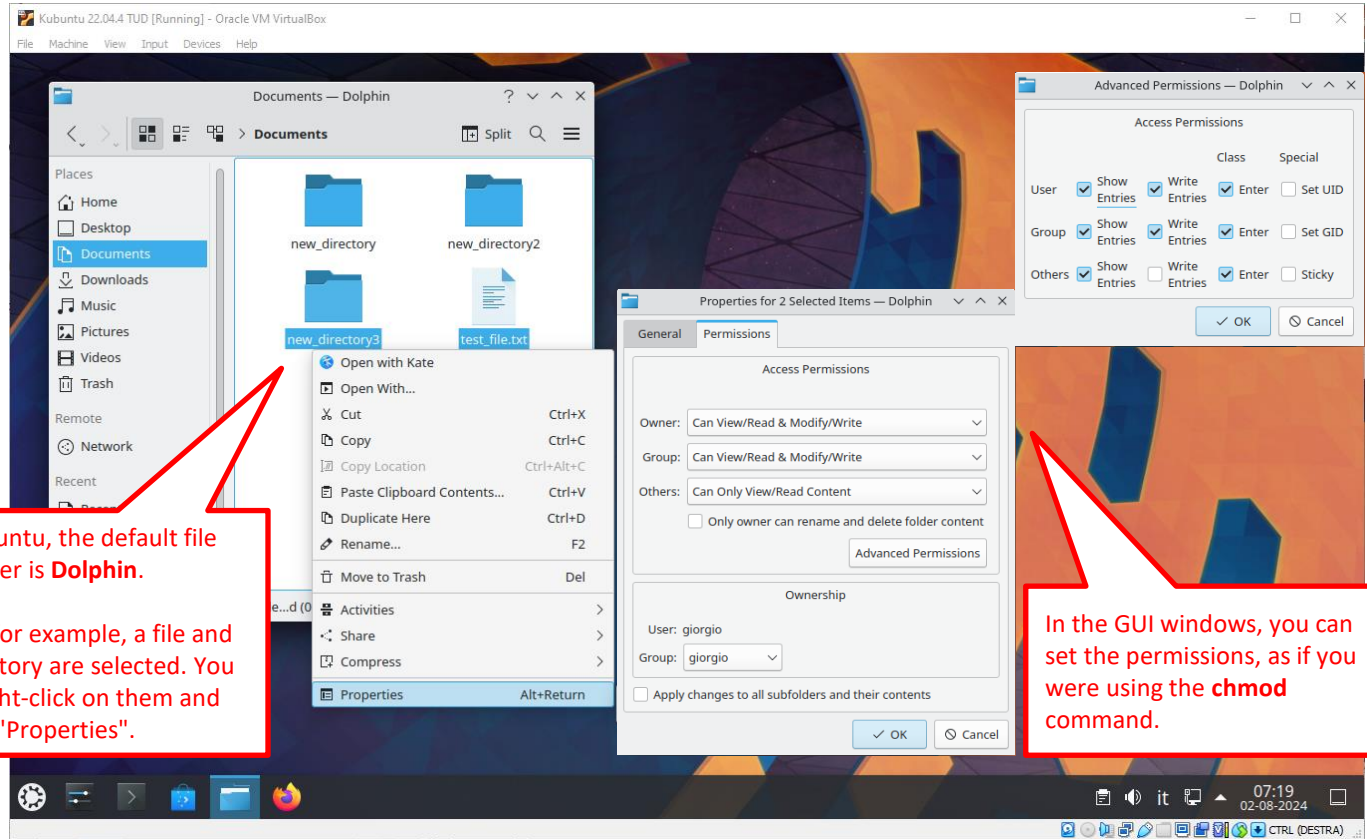
Examples:

- `chmod u+wx filename`: add write and execute privileges to filename for the owner
- `chmod a-w filename`: remove write privilege from filename for everybody
- `chmod ug=rwx filename`: set the privileges of filename to be read, write and execute only for the owner and the owner's group.

File system permissions

Of course, you can also change the permissions of files and directories using the GUI

Linux in a nutshell
Terminal & shell(s)
File system
File system permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
Bash shell scripting
Further resources



In Kubuntu, the default file manager is **Dolphin**.

Here, for example, a file and a directory are selected. You can right-click on them and select "Properties".

In the GUI windows, you can set the permissions, as if you were using the **chmod** command.

Index

- Linux in a nutshell
- The terminal and the shell(s)
- File system
- File system permissions
- **Data streams**
- Processes
- Scheduling jobs
- Installation of software applications
- Text editors
- Bash shell scripting
- Further resources

Data streams

A data stream is, as the name says, a stream of data —especially text data— being passed from one file, device, or program to another.

The GNU Utilities, the Linux core utilities, and many other command-line tools exchange data and perform their work based on data streams.

In Linux and other Unix-like OSes, the use of **Standard Input/Output (STDIO)** is a fundamental way to exchange data between programs: Programs implementing STDIO use standardised file handles for input and output instead of files stored on a disk (or elsewhere).

STDIO is a **buffered data stream**, and its function is to stream data from the output of one program/file/device to the input of another program/file/device.

Linux in a nutshell
Terminal & shell(s)
File system
File system
permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
Bash shell scripting
Further resources

Data streams

There are **3 STDIO data streams**:

- **STDIN** (File handle 0) is the **standard input** which is usually input from the keyboard. STDIN can be redirected from any file, including device files, instead of the keyboard.
- **STDOUT** (File handle 1) is the **standard output** which sends the data stream to the display by default. It is common to redirect STDOUT to a file or to pipe it to another program for further processing.
- **STDERR** (File handle 2) is the **standard error** data stream, i.e. where the program sends error and diagnostics messages. STDERR is also usually sent to the display. If STDOUT is redirected to a file, STDERR continues to be displayed on the screen. STDERR can also be redirected to the same or passed on to the next transformer program in a pipeline.

Data streams: Redirection

Linux includes redirection commands for each stream. These can be used to write STDIN, STDOUT and STDERR to a file. If you write to a file that does not exist, a new file with that name will be created prior to writing.

Commands with a **single bracket overwrite** the destination's existing contents.

> : send to standard output

< : read from standard input

2> : send to standard error

Commands with a **double bracket append** (do not overwrite) the destination's existing contents:

>> append to standard output

<< read from standard input, line by line

2>> append to standard error

Linux in a nutshell
Terminal & shell(s)

File system

File system

permissions

Data streams

Processes

Scheduling jobs

Software install

Text editors

Bash shell scripting

Further resources

Data streams: Redirection

Examples:

`ls /home/giorgio/Documents > giorgio_documents.txt` : List all files/directories in the given directory and write the results to a file

`mkdir " 2> error.txt` : creating a directory with an empty name is not permitted. The error message will be written to file error.txt

`echo Write this message to a new file > file.txt`

`echo Append this line to an existing file >> file.txt` : Example to highlight the difference between overwriting and appending when using `>` and `>>` redirectors

Assuming we have a script called "script.sh":

`script.sh < input_file` : run the script and read from input_file

`script.sh > output_file` : run the script and write to output_file

`script.sh < input_file > output_file` : run the script, read from input_file and write to output_file

`script.sh 2> error_file` : run the script and write the STDERR to error_file

`script.sh > all_output_file 2>&1` : run the script and write both STDOUT and STDERR to all_output_file. **Please note:** In Bash, the last command can be also written as `script.sh >& all_output_file` or `script.sh &> all_output_file`

`script.sh > /dev/null` : Discard the STDOUT, i.e. redirect it to the special device /dev/null

`script.sh 2> /dev/null` : Discard the STDERR, i.e. redirect it to the special device /dev/null

Data streams: Pipes

Pipes are a particular type of stream redirection. They are used to redirect a stream from one program directly to another. When a program's STIOOUT is sent to another through a pipe, the first program's output will be used as STIN to the second, rather than being printed to the terminal. Only the data returned by the second program will be displayed.

In Linux, and other Unix-like OSes, the pipe is represented by a **vertical bar** "|"

The general syntax is:

```
command_1 | command_2 | command_3 | ... | command_N
```

Pipes are unidirectional i.e., **data flows from left to right through the pipeline.**

More details: [https://en.wikipedia.org/wiki/Pipeline_\(Unix\)](https://en.wikipedia.org/wiki/Pipeline_(Unix))

Data streams: Filters

Filters are a class of programs that take plain text as standard input, transforms it into a meaningful format, and then returns it as standard output. They are commonly used with output piped from another program.

cut: extract sections from each line of input, usually from a file

find: returns files with filenames that match the argument passed to find

grep: returns text that matches the string pattern passed to grep

sort: prints the lines of its input in sorted order

tee: redirects standard input to both standard output and one or more files

tr: finds-and-replaces one string with another

uniq: outputs the text with adjacent identical lines collapsed to one, unique line of text

wc ("word count"): counts characters, lines, and words

More details: [https://en.wikipedia.org/wiki/Filter_\(software\)](https://en.wikipedia.org/wiki/Filter_(software))

Data streams: Filters

Examples:

`ls | grep image1.jpg` : List all files and directories, pass them to grep to search for image1.jpg

`ls -l sort` : List all files and directories and return them sorted

`cat long_text.txt | head -15 | tail -5` : selects the first 15 lines, from which the last 5 lines will be eventually displayed

`cat energy_result.csv | grep "PV production" | tee ned.txt` : reads the content of the energy_result file, send it to grep to search for the "PV production" values and finally, with tee, display them in the terminal AND save them to file ned.txt

Chaining operators

The data stream operators are part of a larger family: the **chaining operators**. They are used chain multiple commands together. They are:

| (pipe) : The output of the first command acts as input to the second command (see previous slides)

>, >>, < (redirection) : Redirects the output of a command (see previous slides)

\ (concatenation) : Allows to concatenate long commands spanning over several lines in the shell

() (precedence) : Allows to define the precedence order to execute commands

& (ampersand) : Run a process/script/command in the background

&& (logical "and") : The command following && is only executed if the command preceding && has been successfully executed

|| (logical "or") : The command succeeding || is only executed if the command preceding || has failed.

&&-|| (and-or) : Combination of the && and || operators, similar to the if-then-else statement.

! (not) : Negates an expression within a command

; (semi-colon) : The command following ; is executed even if the command preceding ; has failed

{ } (combination) : The execution of the command list inside { } depends on the execution of the first command in the list

Chaining operators

Examples:

`ping -c20 3d.bk.tudelft.nl &` : Ping the webpage for 20 times, run the process in the background (test it also without **&** to see the difference: can you interact with the terminal?)

`who ; pwd ; ls` : Simply run the 3 commands one after the other, no matter if one fails or succeeds

`echo "Print this!" && echo -e "\nThe first echo command succeeded"` : The second echo is run only if the first one succeeds

`mkdir "" || echo -e "\nThe first command failed"` : the echo is run only because the first command fails (you cannot create a nameless directory)

`ping -c1 3d.bk.tudelft.nl && echo "Successful ping" || echo "Failed ping"` : ping the URL, if there is an answer then execute the first echo, else execute the second echo

`rm !(*.jpg)` : remove (delete) all files that do not have a .jpg extension

`ping -c1 www.brickset.com && { echo -e "\n\n*** The webpage exists!" ; firefox www.brickset.com & }` : Ping the webpage. Then consider the list of commands inside `{ }`. Run the second in the list only if the first (echo) succeeds. Beware the syntax: please note the empty spaces after the `{` and before the `}` parentheses

Index

- Linux in a nutshell
- The terminal and the shell(s)
- File system
- File system permissions
- Data streams
- **Processes**
- Scheduling jobs
- Installation of software applications
- Text editors
- Bash shell scripting
- Further resources

A running instance of a program is called a **process**

- **Example:** If you have opened two terminal windows, most likely you have two processes of the same program (e.g. "Konsole"). Each terminal loads a shell program (e.g. the Bash shell): each running shell is itself another process. Whenever you issue a command from the shell (e.g. "cp"), the corresponding program is executed in a new process, too.

Processes in Linux are organized as a tree.

- The *init* process is the root process
- Each process has its own ID (PID, process ID)
- Each process has the ID of the parent process (PPID, parent process ID)

Multiple processes, running in parallel or in series, can be grouped in jobs. A **job** is a scheduled process or set of processes.

In Linux, there are **5 types** of processes.

- **Parent process:** The process created by the user on the terminal. All processes have a parent process. If it was created directly by the user, then the parent process will be the kernel process
- **Child process:** A child process is a process that is created by another process known as the parent process
- **Orphan process:** A child process becomes an orphan process when the parent gets executed before its own child process. In such a case, the orphan process has a “Init” process (PID 0) as its PPID
- **Zombie process:** A process that is already dead but shows up in process status. Zombie processes have zero CPU consumption
- **Daemon process:** A system-related background processes. These processes often run the permissions of root and service requests from other processes. A Daemon process often runs in the background. A Daemon process can be recognized if it has “?” in its TTY field (see later)

Linux in a nutshell
Terminal & shell(s)

File system

File system

permissions

Data streams

Processes

Scheduling jobs

Software install

Text editors

Bash shell scripting

Further resources

Linux processes can be run in foreground or in background

- **Foreground processes** are started by the user and are the default. They accept command-line input and output it to the computer screen. A running foreground process prevents the start or execution of other, following processes because the command prompt will not be available until the currently running program completes its processing and comes out.
- **Background processes** run, as the name says, in the background. They are non-interactive and do not need keyboard input. While one process is running in the background, it is possible to start another process from the terminal. By adding a **single ampersand ("&")** at the end of a command, the command can be executed as a background process

Example:

- `sleep 5 && echo "Ciao!"`: Force the terminal to wait for 5 seconds, then print the message to the screen. Nothing can be done in-between in the terminal.
- `sleep 10 && echo "Ciao!" &`: Note the last, single **&**! After issuing the command, you can still interact with the terminal. After 5 seconds, the message will be printed to the screen.

Processes

Useful BASH commands for processes

`ps` ("process status"): list processes, with different levels of information, e.g. `ps -fu`

`ps <pid>`: check status of a single process, identified by its PID

`pstree`: print the tree of the processes

`kill <pid>`: end/terminate a process, identified by its PID

`top`: display Linux processes

`htop`: an interactive system control, process viewer, and process manager

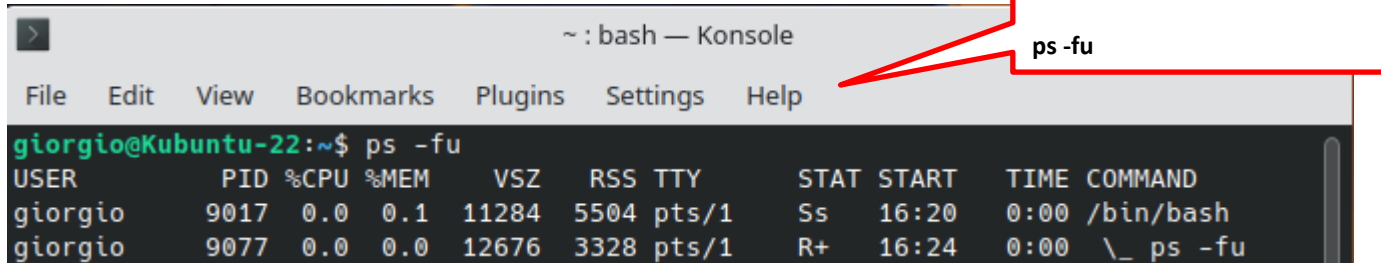
`jobs`: display status of jobs

Miscellaneous commands:

`free`: display the total amount of free and used memory (RAM) on the Linux system

`df`: display the free disk space(Hard Disk) on all the file systems

Processes



```

~ : bash — Konsole
File Edit View Bookmarks Plugins Settings Help
giorgio@Kubuntu-22:~$ ps -fu
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
giorgio   9017  0.0  0.1  11284  5504 pts/1    Ss   16:20   0:00 /bin/bash
giorgio   9077  0.0  0.0  12676  3328 pts/1    R+   16:24   0:00 \_ ps -fu
  
```

USER: Process owner

PID: The process ID

%CPU, %MEM: Percentage of CPU/memory used

VZS: Virtual memory size, virtual memory used by the process (in kB)

RSS: Resident set size, physical memory used by the process (in kB)

TTY: The terminal associated with the process

STAT: The state code of the process; there are many values, but the common ones are S (sleeping) and R (running)

START: The time when the process started

TIME: The CPU time consumed by the process

COMMAND: Command issued that created the process

More details: [https://it.wikipedia.org/wiki/Ps_\(Unix\)](https://it.wikipedia.org/wiki/Ps_(Unix))

- Linux in a nutshell
- Terminal & shell(s)
- File system
- File system
- permissions
- Data streams
- Processes**
- Scheduling jobs
- Software install
- Text editors
- Bash shell scripting
- Further resources

Processes

htop is interactive system control, process viewer, and process manager

- Linux in a nutshell
- Terminal & shell(s)
- File system
- File system permissions
- Data streams
- Processes**
- Scheduling jobs
- Software install
- Text editors
- Bash shell scripting
- Further resources

```

~ : htop — Konsole
File Edit View Bookmarks Plugins Settings Help

0[|||] 3.3%] Tasks: 91, 179 thr; 1 running
1[||] 2.0%] Load average: 0.06 0.09 0.13
Mem[|||||||||||||||||||||||||||||||||||||||||]876M/3.82G Uptime: 1 day, 03:11:32
Swp[|] 524K/3.81G

  PID USER   PRI  NI  VIRT   RES   SHR  S  CPU% MEM%  TIME+  Command
 950 root     20   0  338M  113M  77320 S  1.3  2.9  1:44.12 /usr/lib/xorg/Xorg -nolisten tcp -auth /var/run/sddm/{7e108366-4d95-455d-b
1219 giorgio 20   0  213M  3100  2688 S  0.0  0.1  0:10.20 /usr/bin/VBoxClient --seamless
1233 giorgio 20   0  213M  3460  3072 S  0.0  0.1  0:09.85 /usr/bin/VBoxClient --vmsvga-session
1351 giorgio 20   0  1200M  195M  137M S  2.0  5.0  2:11.40 /usr/bin/kwin_x11
1357 giorgio 20   0  1200M  195M  137M S  0.0  5.0  0:09.90 /usr/bin/kwin_x11
1370 giorgio 20   0  1200M  195M  137M S  0.0  5.0  0:37.70 /usr/bin/kwin_x11
1397 giorgio 20   0  1911M  340M  163M S  2.0  8.7  0:46.08 /usr/bin/plasmashell
9251 giorgio 20   0  601M  119M  99796 S  0.0  3.1  0:01.31 /usr/bin/konsole -qwindowtitle Htop -qwindowicon htop -e /usr/bin/htop
9258 giorgio 20   0  11232  4864  3584 R  0.7  0.1  0:00.19 /usr/bin/htop
  1 root     20   0  163M  13012  8276 S  0.0  0.3  0:07.14 /sbin/init splash
 254 root     19  -1  56464  25092  23940 S  0.0  0.6  0:00.49 /lib/systemd/systemd-journald
 283 root     20   0  26844  6912  4608 S  0.0  0.2  0:00.32 /lib/systemd/systemd-udev
 432 systemd-r 20   0  25804  13536  9088 S  0.0  0.3  0:00.27 /lib/systemd/systemd-resolved
 448 root     20   0  8584  4872  1792 S  0.0  0.1  0:01.58 /usr/sbin/haveged --Foreground --verbose=1
 466 root     20   0  234M  7808  6912 S  0.0  0.2  0:01.13 /usr/libexec/accounts-daemon
 467 root     20   0  2816  1920  1792 S  0.0  0.0  0:00.30 /usr/sbin/acpid
 470 avahi    20   0  7632  4096  3712 S  0.0  0.1  0:00.03 avahi-daemon: running [Kubuntu-22.local]
 471 root     20   0  9500  2944  2688 S  0.0  0.1  0:00.09 /usr/sbin/cron -f -P
 472 messagebu 20   0  10276  6272  4096 S  0.0  0.2  0:07.91 @dbus-daemon --system --address=systemd: --nofork --nopidfile --systemd-ac
 473 root     20   0  254M  18176  15360 S  0.0  0.5  0:01.37 /usr/sbin/NetworkManager --no-daemon
 484 root     20   0  82700  3840  3584 S  0.0  0.1  0:02.07 /usr/sbin/irqbalance --foreground
 491 root     20   0  35440  19572  10624 S  0.0  0.5  0:00.11 /usr/bin/python3 /usr/bin/networkd-dispatcher --run-startup-triggers
 494 root     20   0  237M  11324  7424 S  0.0  0.3  0:00.62 /usr/libexec/polkitd --no-debug
 496 root     20   0  82700  3840  3584 S  0.0  0.1  0:00.00 /usr/sbin/irqbalance --foreground
 497 root     20   0  234M  7552  6784 S  0.0  0.2  0:00.02 /usr/libexec/power-profiles-daemon
 498 syslog   20   0  217M  6016  4480 S  0.0  0.2  0:00.08 /usr/sbin/rsyslogd -n -iNONE
 501 root     20   0  237M  11324  7424 S  0.0  0.3  0:00.00 /usr/libexec/polkitd --no-debug

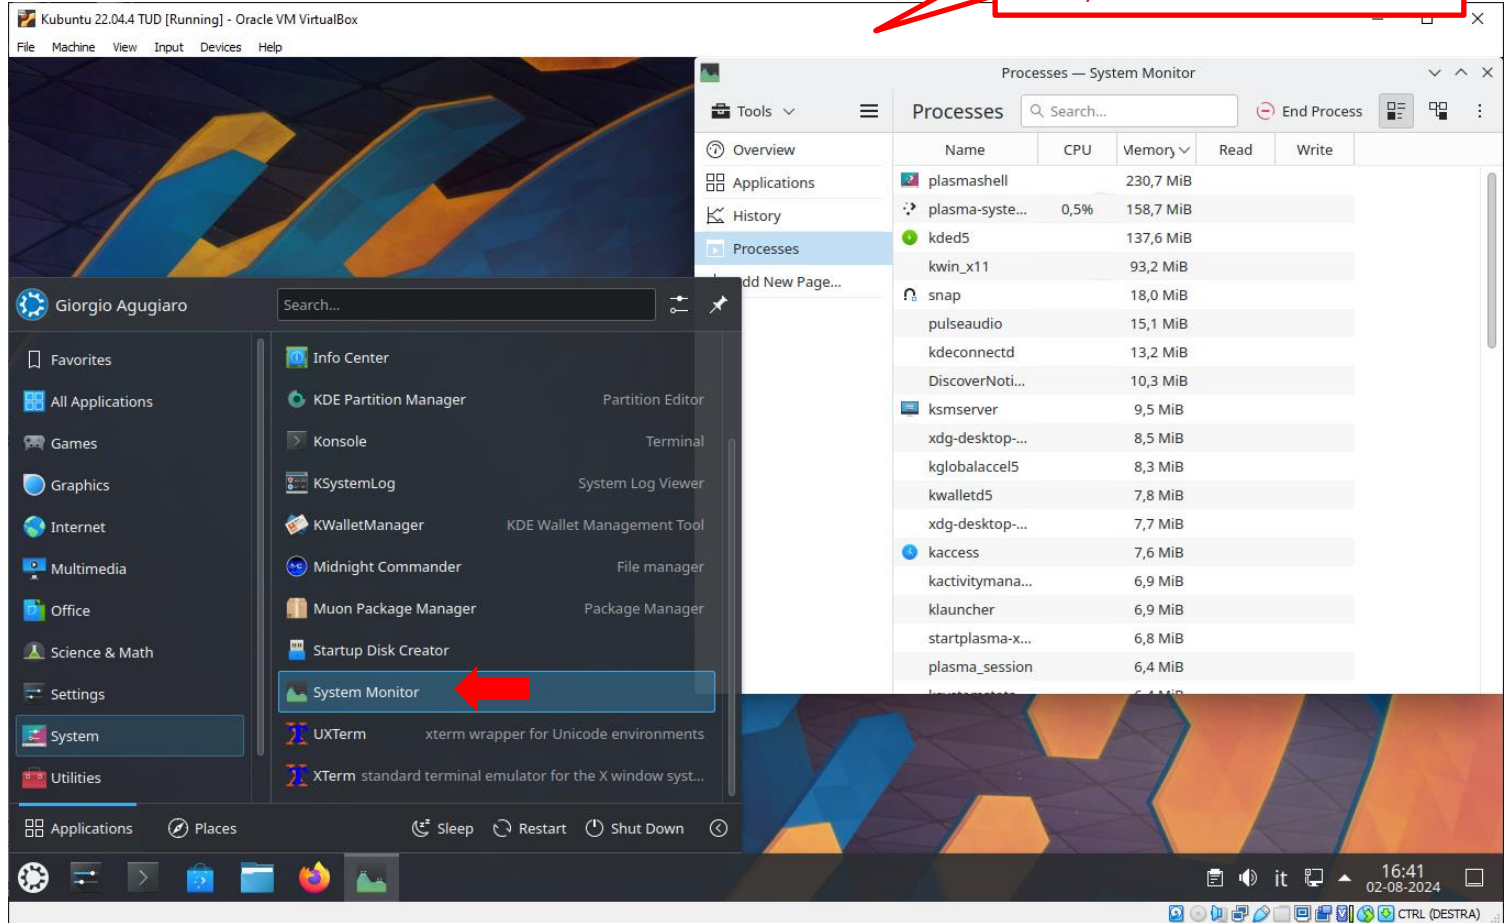
F1Help F2Setup F3Search F4Filter F5Tree F6SortBy F7Nice +F8Nice +F9Kill F10Quit

```

Processes

- Linux in a nutshell
- Terminal & shell(s)
- File system
- File system
- permissions
- Data streams
- Processes**
- Scheduling jobs
- Software install
- Text editors
- Bash shell scripting
- Further resources

Finally, you also can use the "System monitor" to manage the processes directly from the Kubuntu GUI



Kubuntu 22.04.4 TUD [Running] - Oracle VM VirtualBox

File Machine View Input Devices Help

Processes — System Monitor

Tools Overview Applications History Processes

| Name | CPU | Memory | Read | Write |
|------------------|------|-----------|------|-------|
| plasmashell | | 230,7 MiB | | |
| plasma-syste... | 0,5% | 158,7 MiB | | |
| kded5 | | 137,6 MiB | | |
| kwin_x11 | | 93,2 MiB | | |
| snap | | 18,0 MiB | | |
| pulseaudio | | 15,1 MiB | | |
| kdeconnectd | | 13,2 MiB | | |
| DiscoverNoti... | | 10,3 MiB | | |
| kmsserver | | 9,5 MiB | | |
| xdg-desktop-... | | 8,5 MiB | | |
| kglobalaccl5 | | 8,3 MiB | | |
| kwalletd5 | | 7,8 MiB | | |
| xdg-desktop-... | | 7,7 MiB | | |
| kaccess | | 7,6 MiB | | |
| kactivitymana... | | 6,9 MiB | | |
| klauncher | | 6,9 MiB | | |
| startplasma-x... | | 6,8 MiB | | |
| plasma_session | | 6,4 MiB | | |
| ... | | 6,4 MiB | | |

Giorgio Agugiaro Search...

Info Center KDE Partition Manager Konsole KSystemLog KWalletManager Midnight Commander Muon Package Manager Startup Disk Creator System Monitor UXTerm XTerm

Applications Places Sleep Restart Shut Down

16:41 02-08-2024 CTRL (DESTRA)

Index

- Linux in a nutshell
- The terminal and the shell(s)
- File system
- File system permissions
- Data streams
- Processes
- **Scheduling jobs**
- Installation of software applications
- Text editors
- Bash shell scripting
- Further resources

Scheduling jobs

A process or a group of processes (a job) can be scheduled to be run at a certain point in time, or with a certain schedule. Relevant commands are:

at: execute commands for one time at a specified time (in future)

atq (at queue): print the list of user's pending jobs

cron: a Daemon process. It reads every minute the crontab table and executes the scheduled jobs

crontab: show and manage the table containing the list of scheduled jobs. In particular

- **crontab -l**: list all jobs in the crontab table
- **crontab -e**: open the editor to add, remove, change the scheduled jobs in the crontab table
- **crontab -r**: remove the complete list of scheduled jobs

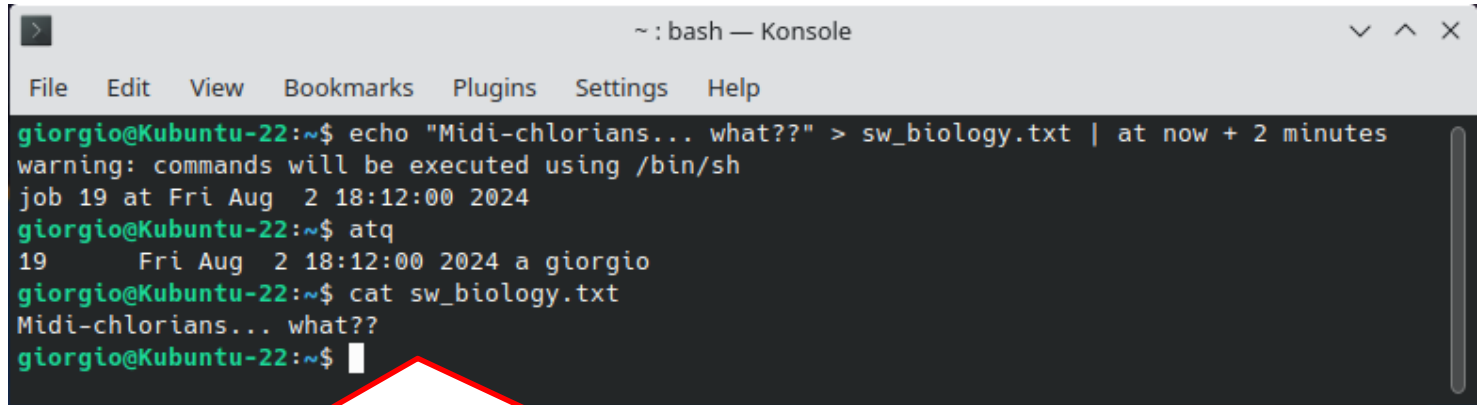
Please note: For **at** and **cron** there are several command options and some ancillary commands. Please refer to the documentation for further details. The next slides will provide only some simple examples

at: [https://en.wikipedia.org/wiki/At_\(Unix\)](https://en.wikipedia.org/wiki/At_(Unix))

cron: <https://en.wikipedia.org/wiki/Cron>

Scheduling jobs: at

The command **at** can be run in different ways: **Example 1**



```
~ : bash — Konsole
File Edit View Bookmarks Plugins Settings Help
giorgio@Kubuntu-22:~$ echo "Midi-chlorians... what??" > sw_biology.txt | at now + 2 minutes
warning: commands will be executed using /bin/sh
job 19 at Fri Aug 2 18:12:00 2024
giorgio@Kubuntu-22:~$ atq
19      Fri Aug 2 18:12:00 2024 a giorgio
giorgio@Kubuntu-22:~$ cat sw_biology.txt
Midi-chlorians... what??
giorgio@Kubuntu-22:~$
```

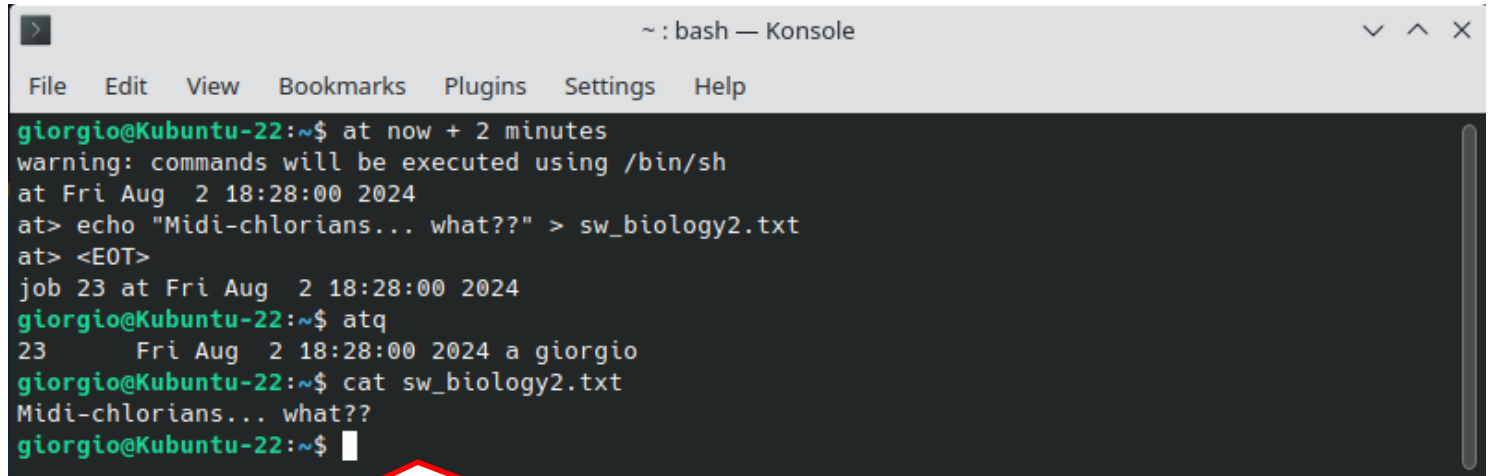
Using piping, you echo a message that will be written to a text file. This operation is to be carried out in 2 minutes from now. A job is therefore created.

You can check the pending job with **atq**.

After 2 minutes, you can check that the file has been written and contains the original message.

Scheduling jobs: at

The command **at** can be run in different ways: **Example 2**



```
~ : bash — Konsole
File Edit View Bookmarks Plugins Settings Help
giorgio@Kubuntu-22:~$ at now + 2 minutes
warning: commands will be executed using /bin/sh
at Fri Aug 2 18:28:00 2024
at> echo "Midi-chlorians... what??" > sw_biology2.txt
at> <EOT>
job 23 at Fri Aug 2 18:28:00 2024
giorgio@Kubuntu-22:~$ atq
23      Fri Aug 2 18:28:00 2024 a giorgio
giorgio@Kubuntu-22:~$ cat sw_biology2.txt
Midi-chlorians... what??
giorgio@Kubuntu-22:~$
```

Here you perform the same operation as before, however using **at's interactive prompt** that allows you to enter which commands to run at the specified time. A warning stating which shell the command will use is also printed.

You can exit the interactive prompt and save the scheduled job by pressing **Ctrl + D**. You can cancel the job with **Ctrl + C**.

Scheduling jobs: cron

Scheduling jobs for **cron** consists in adding lines to the crontab table. Assume that you have a shell script named "script.sh" that you want to run at regular intervals. The generic syntax to add a job entry to crontab is

```
A B C D E COMMAND
```

with:

A: Minutes range from 0 to 59; default is * (i.e. all values)

B: Hours range from 0 to 23; default is * (i.e. all values)

C: Days range from 0 to 31; default is * (i.e. all values)

D: Months range from 0 to 12; default is * (i.e. all values)

E: Days of the week range from 0 to 7 (Sunday is 0 or 7); default is * (i.e. all values)

COMMAND: command to be executed.

Example:

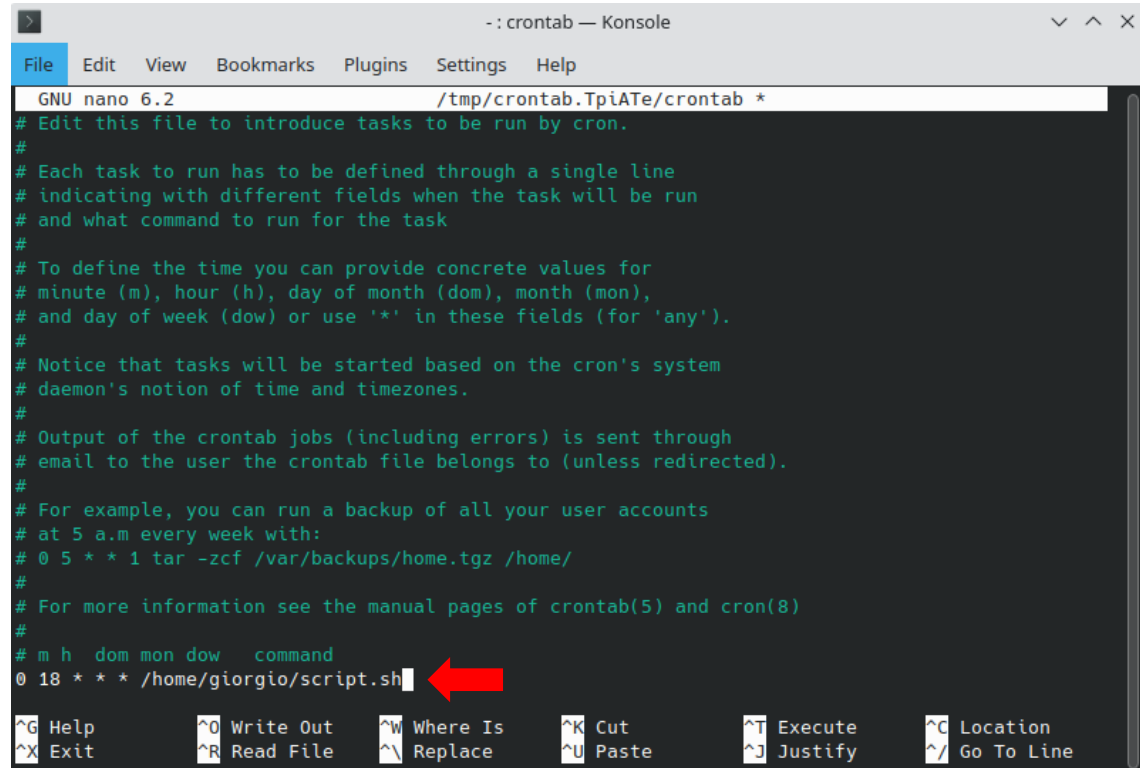
Line to add to run the script every day of the year exactly at 18:00 (6:00 PM)

```
0 18 * * * /home/giorgio/script.sh
```

Scheduling jobs: cron

Run `crontab -e` to edit the crontab table and add a line for each job at the end of the file. Save and exit. Suggestion: always use ABSOLUTE PATHS to your script.

Linux in a nutshell
 Terminal & shell(s)
 File system
 File system
 permissions
 Data streams
 Processes
Scheduling jobs
 Software install
 Text editors
 Bash shell scripting
 Further resources

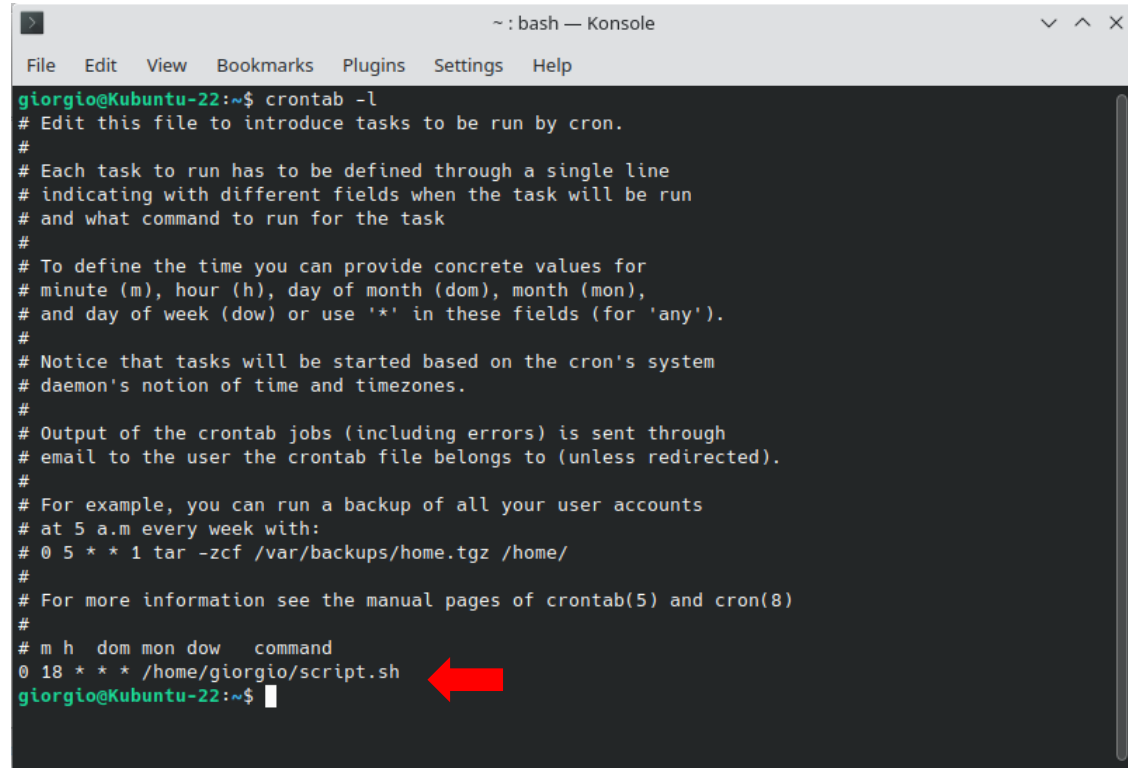


```

- : crontab — Konsole
File Edit View Bookmarks Plugins Settings Help
GNU nano 6.2 /tmp/crontab.TpiATe/crontab *
# Edit this file to introduce tasks to be run by cron.
#
# Each task to run has to be defined through a single line
# indicating with different fields when the task will be run
# and what command to run for the task
#
# To define the time you can provide concrete values for
# minute (m), hour (h), day of month (dom), month (mon),
# and day of week (dow) or use '*' in these fields (for 'any').
#
# Notice that tasks will be started based on the cron's system
# daemon's notion of time and timezones.
#
# Output of the crontab jobs (including errors) is sent through
# email to the user the crontab file belongs to (unless redirected).
#
# For example, you can run a backup of all your user accounts
# at 5 a.m every week with:
# 0 5 * * 1 tar -zcf /var/backups/home.tgz /home/
#
# For more information see the manual pages of crontab(5) and cron(8)
#
# m h dom mon dow  command
0 18 * * * /home/giorgio/script.sh
  
```

Scheduling jobs: cron

Run `crontab -l` to print to screen the crontab table. Here you can see that the job has been added and will be carried out at the set time.



```
~ : bash — Konsole
File Edit View Bookmarks Plugins Settings Help
giorgio@Kubuntu-22:~$ crontab -l
# Edit this file to introduce tasks to be run by cron.
#
# Each task to run has to be defined through a single line
# indicating with different fields when the task will be run
# and what command to run for the task
#
# To define the time you can provide concrete values for
# minute (m), hour (h), day of month (dom), month (mon),
# and day of week (dow) or use '*' in these fields (for 'any').
#
# Notice that tasks will be started based on the cron's system
# daemon's notion of time and timezones.
#
# Output of the crontab jobs (including errors) is sent through
# email to the user the crontab file belongs to (unless redirected).
#
# For example, you can run a backup of all your user accounts
# at 5 a.m every week with:
# 0 5 * * 1 tar -zcf /var/backups/home.tgz /home/
#
# For more information see the manual pages of crontab(5) and cron(8)
#
# m h dom mon dow  command
0 18 * * * /home/giorgio/script.sh
giorgio@Kubuntu-22:~$
```

Linux in a nutshell
Terminal & shell(s)
File system
File system
permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
Bash shell scripting
Further resources

Index

- Linux in a nutshell
- The terminal and the shell(s)
- File system
- File system permissions
- Data streams
- Processes
- Scheduling jobs
- **Installation of software applications**
- Text editors
- Bash shell scripting
- Further resources

Installation of software applications

- In Linux there are several ways to install a software application
- The easiest, and by far the most common way, is to install software applications from a **repository**. A repository is a public server that hosts software packages
- A **software package** (or just: package) is a collection of files and metadata that contains a specific software application. The purpose is to simplify the process of distributing, installing, and managing software on a computer system
- There are different types of software **package formats**, depending on the distribution. The two most common ones are **rpm** and **deb**. The *rpm* format is tailored to Red Hat Linux and its derivatives, while the *deb* format is for Debian-based distributions, such as (K)Ubuntu
- A Linux distribution provides a command, and usually a GUI-based program, that retrieves the software from a repository and installs it onto your computer
 - It is conceptually the same as searching for and installing an app on your smartphone from the Android or Apple stores!

Linux in a nutshell
Terminal & shell(s)
File system
File system
permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
Bash shell scripting
Further resources

Installation from a repository

- The **installation from a repository** can be carried out, as usual, in two ways:
 - a) Using the terminal
 - b) Using a GUI based software installation program
- When using the terminal, you only need to know the exact name of the package containing the program you want to install
- Before installing the software itself, it is a good habit to refresh the database(s) on your machine that contains the list of applications available in the repositories
- All these operations require superuser privileges, but can be carried out also by a normal user thanks to **sudo**
 - In (K)Ubuntu, you use the **apt** command to perform software installation and other related operations
- Alternatively, you can use the GUI-based package manager to search for the package name, and to install or remove it without using the terminal at all
 - In Kubuntu, the GUI-based package manager is called **Discover**
 - In Ubuntu, the GUI-based package manager is called **Ubuntu Software Center**

Installation from a repository

Example 1:

Let's assume we want to install the file manager **Midnight Commander**. Its package name is **mc**. Let's open a terminal and type (you will be asked your password):

```
sudo apt update
```

```
sudo apt install mc
```

That's all! 😊 The first line updates the database(s) containing information on the software in the repositories. The second line downloads and installs Midnight Commander. At the end, you can type **mc**, and the program will be launched. Using command chaining, the above commands can be also written as

```
sudo apt update && sudo apt install mc
```

Finally, if you want to remove/uninstall a software package (e.g. **mc**), you simply type:

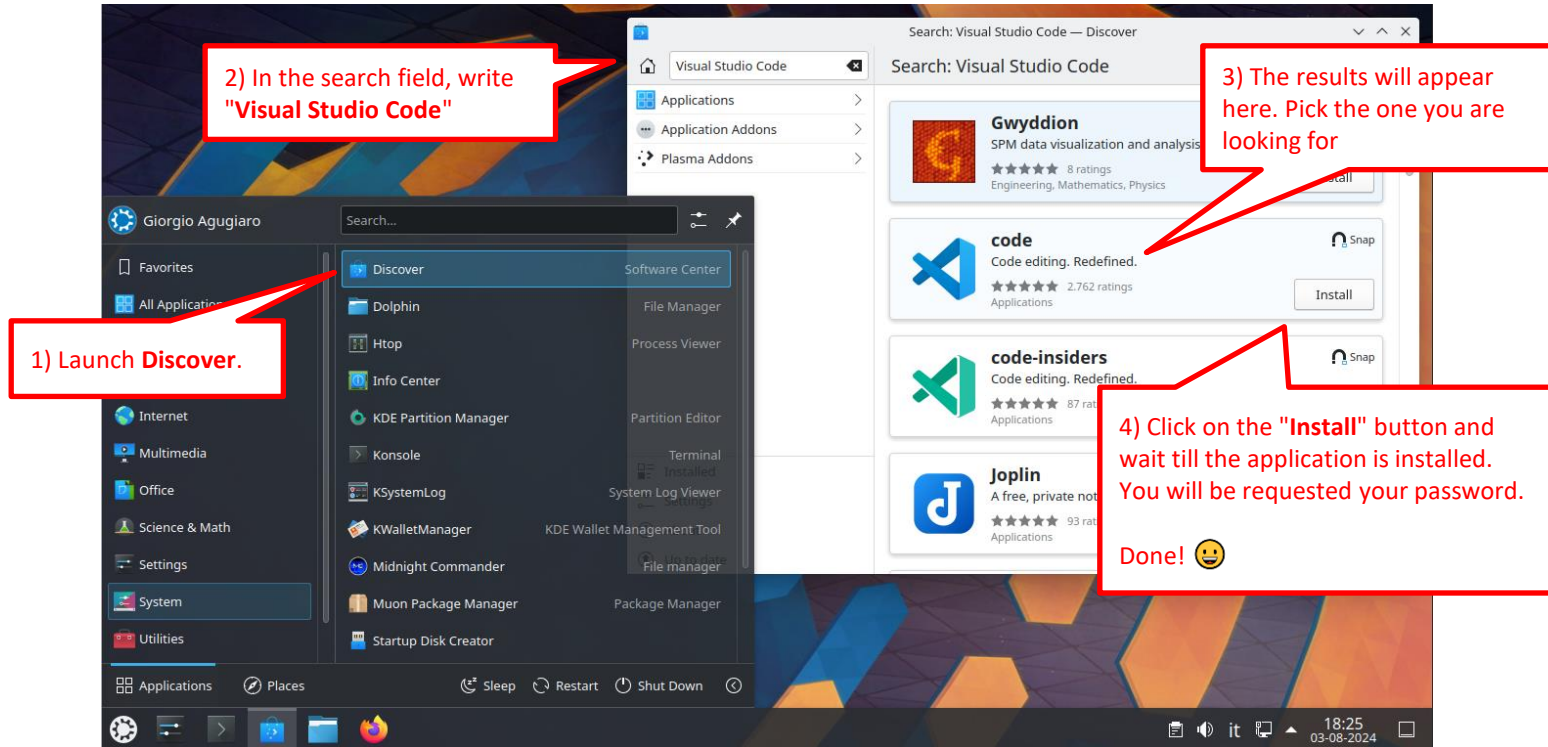
```
sudo apt remove mc
```

Installation from a repository

Example 2:

Let's assume we want to install **Visual Studio Code**. But you do not know the exact package name. Let's use **Discover** in Kubuntu.

Linux in a nutshell
Terminal & shell(s)
File system
File system
permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
Bash shell scripting
Further resources



2) In the search field, write "Visual Studio Code"

3) The results will appear here. Pick the one you are looking for

1) Launch Discover.

4) Click on the "Install" button and wait till the application is installed. You will be requested your password.

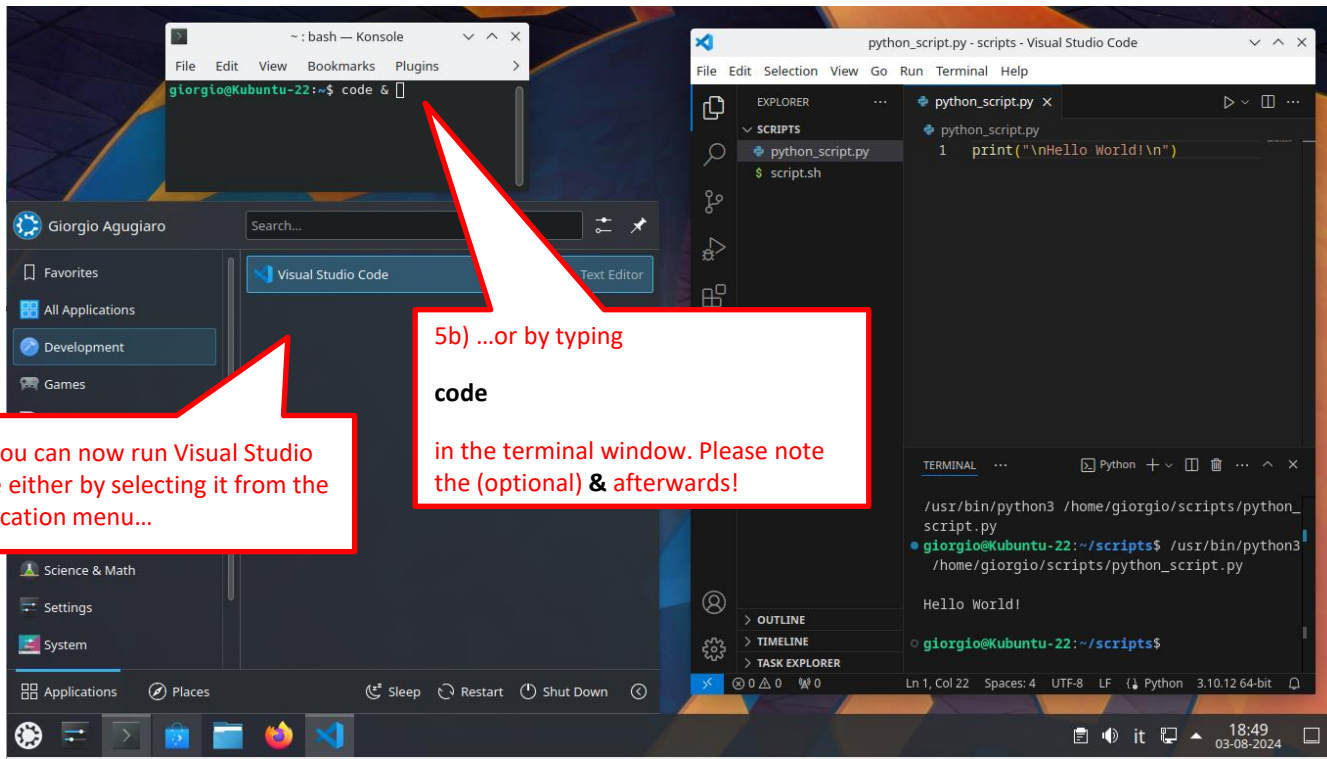
Done! 😊

Installation from a repository

Example 2:

Launch Visual Studio Code

Linux in a nutshell
 Terminal & shell(s)
 File system
 File system
 permissions
 Data streams
 Processes
 Scheduling jobs
Software install
 Text editors
 Bash shell scripting
 Further resources



5a) You can now run Visual Studio Code either by selecting it from the application menu...

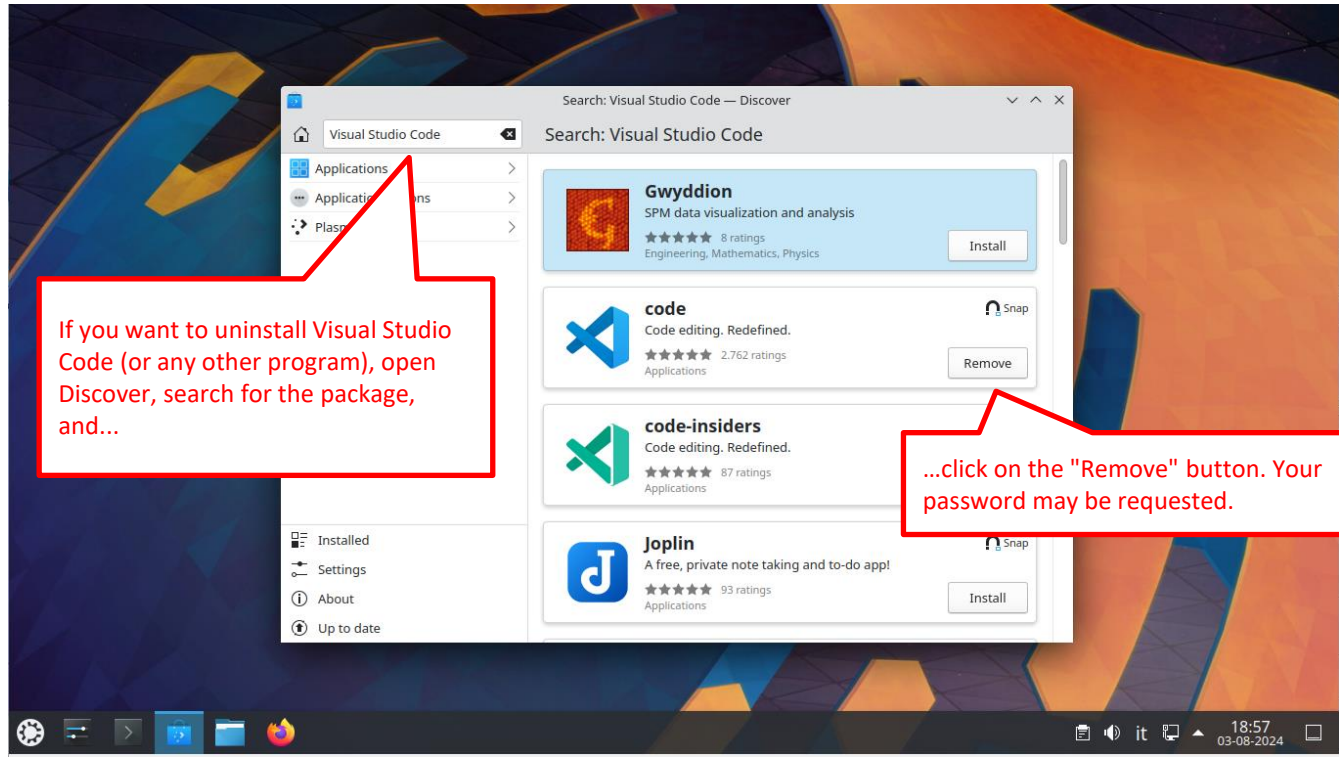
5b) ...or by typing `code` in the terminal window. Please note the (optional) `&` afterwards!

Installation from a repository

Example 2:

Uninstall Visual Studio Code

Linux in a nutshell
Terminal & shell(s)
File system
File system
permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
Bash shell scripting
Further resources



Index

- Linux in a nutshell
- The terminal and the shell(s)
- File system
- File system permissions
- Data streams
- Processes
- Scheduling jobs
- Installation of software applications
- **Text editors**
- Bash shell scripting
- Further resources

Text editors

- A fundamental application that can't miss in a Linux machine is a **text editor**. As a matter of fact, there are several text editors for Linux. Some are based on the command-line, others exploit the GUI possibilities of the Desktop Environment they are part of
- In terms of command-line text editors, the two most well-known are **vim** and **nano**
- **Kate** is the text editor shipped with KDE (e.g. in Kubuntu), while **gedit** is the "equivalent" that comes with GNOME (e.g. in Ubuntu).
- But there are many more, even several extensions for **Visual Studio Code** that provide support for the Bash shell!



Further details: https://en.wikipedia.org/wiki/Category:Unix_text_editors

Text editors: Nano



In the realm of command-line text editors, **nano** (more precisely: GNU nano) is less powerful than **vim**, but much easier to learn and use. It can be started simply typing `nano` or `nano <file_name>`. It provides a two-line shortcut bar at the bottom of the screen which lists the available commands.

Some of these commands are:

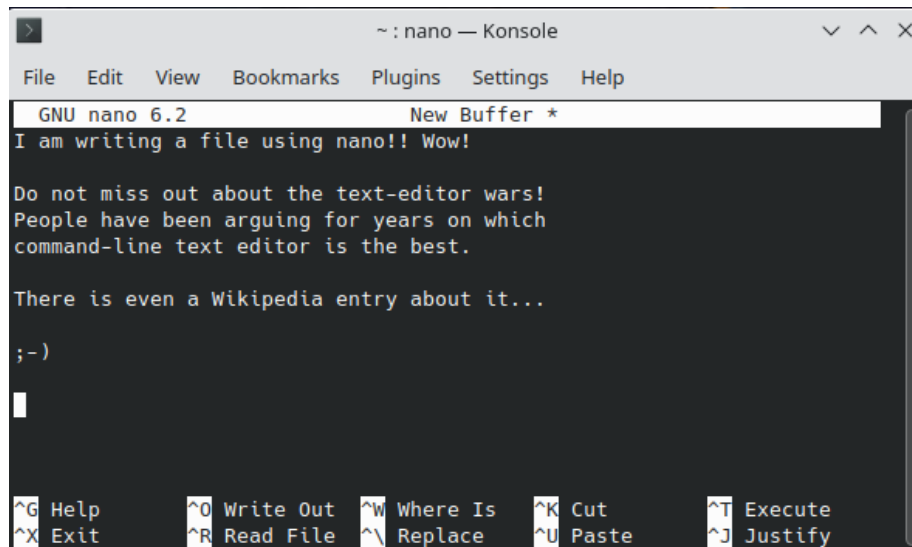
CTRL-R: Insert contents from another file to the current buffer

CTRL-G: Display the help screen

CTRL-O: Write to a file

CTRL-X: Exit a file

CTRL-C: Show cursor position



```

> ~: nano — Konsole
File Edit View Bookmarks Plugins Settings Help
GNU nano 6.2 New Buffer *
I am writing a file using nano!! Wow!


Do not miss out about the text-editor wars!
People have been arguing for years on which
command-line text editor is the best.

There is even a Wikipedia entry about it..

;-)
|

^G Help      ^O Write Out  ^W Where Is   ^K Cut        ^T Execute
^X Exit      ^R Read File  ^\ Replace    ^U Paste      ^J Justify
  
```

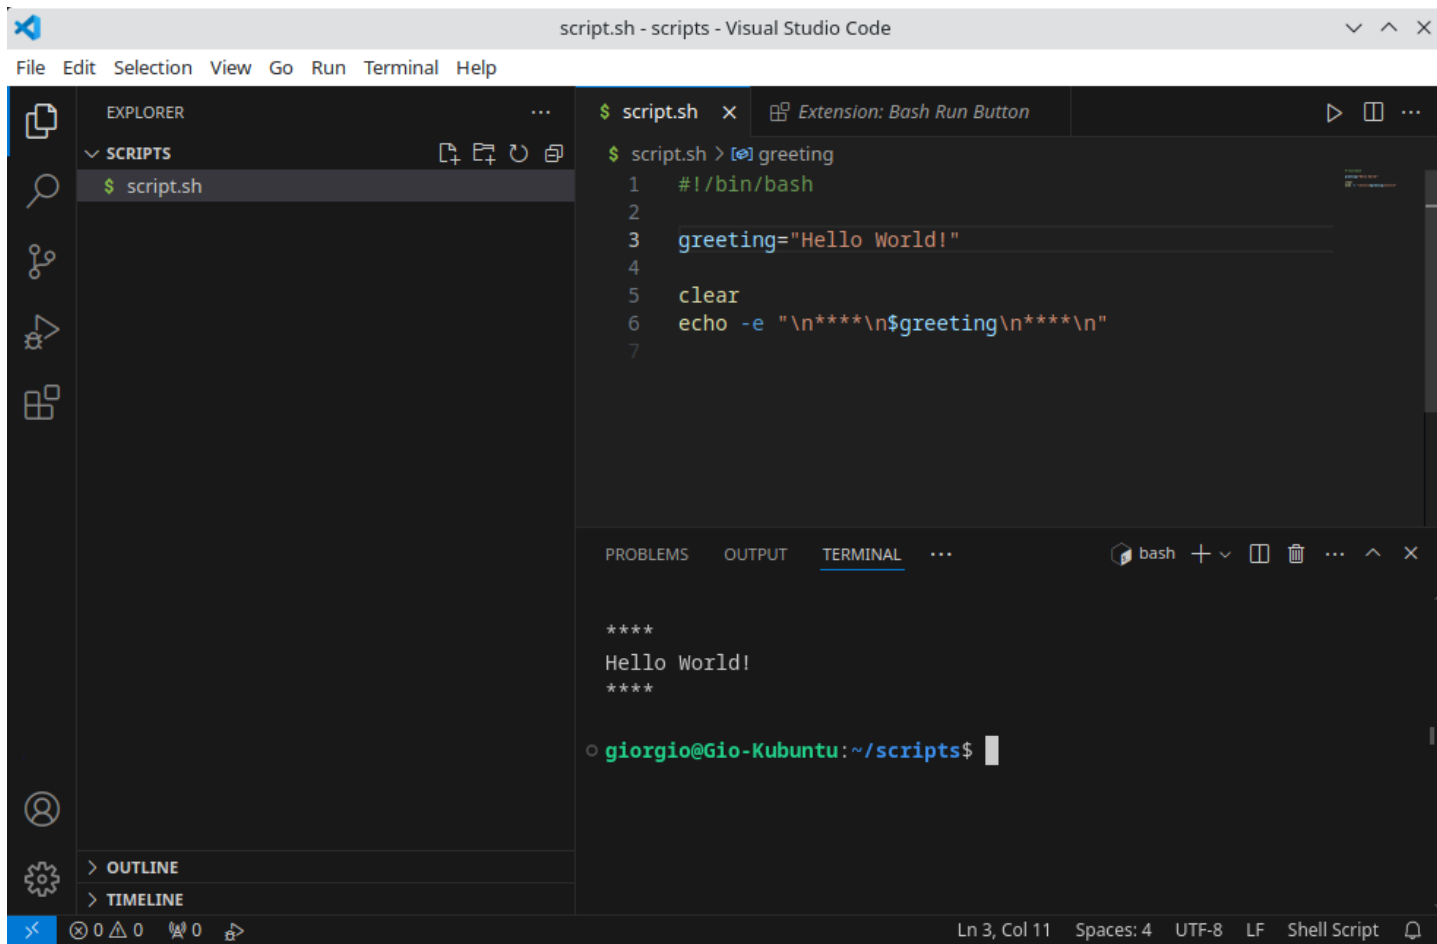
- Linux in a nutshell
- Terminal & shell(s)
- File system
- File system
- permissions
- Data streams
- Processes
- Scheduling jobs
- Software install
- Text editors**
- Bash shell scripting
- Further resources



Text editors: Visual Studio Code



- Linux in a nutshell
- Terminal & shell(s)
- File system
- File system
- permissions
- Data streams
- Processes
- Scheduling jobs
- Software install
- Text editors**
- Bash shell scripting
- Further resources



The screenshot shows the Visual Studio Code interface with a terminal window open. The terminal displays the execution of a shell script named 'script.sh'. The script content is as follows:

```
$ script.sh > greeting
1  #!/bin/bash
2
3  greeting="Hello World!"
4
5  clear
6  echo -e "\n****\n${greeting}\n****\n"
7
```

The terminal output shows the script's execution result:

```
****
Hello World!
****

giorgio@Gio-Kubuntu: ~/scripts$
```

The status bar at the bottom indicates the current position in the file: Ln 3, Col 11, Spaces: 4, UTF-8, LF, Shell Script.

Index

- Linux in a nutshell
- The terminal and the shell(s)
- File system
- File system permissions
- Data streams
- Processes
- Scheduling jobs
- Installation of software applications
- Text editors
- **Bash shell scripting**
- Further resources

Shell scripting

Linux in a nutshell
Terminal & shell(s)
File system
File system
permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
**Bash shell
scripting**
Further resources

- **Shell scripting** consists in turning a series of commands into a script that can be run as many times as needed Shell scripting is primarily meant to automate repetitive tasks, test solutions, and increase efficiency
- Examples of tasks that benefit from shell scripting can be performing backups of files, monitoring system resources, and managing user accounts
- A shell script is in essence a text file with a list of commands that instruct an operating system to perform certain operations
- A shell script must be read and interpreted by a shell program
- A shell script file has generally a `.sh` extension
- A shell script can be run in two ways:
 - As argument of the shell binary/executable file (e.g. `/bin/bash`)
 - Example: `/bin/bash ./my_script.sh`
 - As an executable file, which however must be made executable with `chmod u+x my_script.sh`
 - Example: `./my_script.sh`

- Shell script files must begin with the so-called **shebang**. The first line of the script contains the absolute path to the shell interpreter. This is relevant especially when there exist different shells in the same machine
- The shebang is written as `#!/bin/bash` (for the Bash shell)
- The name comes from the combination of the terms sharp (#) and bang (!). Besides shebang, it is also known as sha-bang, hashbang, etc.

```
1  #!/bin/bash
2
3  echo "May the Bash be with you!"
```

Example of a simple shell script that just prints a message to the screen. The first line is the shebang

More details: [https://en.wikipedia.org/wiki/Shebang_\(Unix\)](https://en.wikipedia.org/wiki/Shebang_(Unix))

Please note: Providing a full guide to shell scripting is beyond the purpose of this introductory guide. Only the basics will be mentioned here. A good starting point is the open-source book "**Introduction to Bash scripting**" by Bobby Iliev, available also on GitHub at <https://github.com/bobbyiliev/introduction-to-bash-scripting>.

Comments:

Comments must be preceded by the **#** symbol. Example:

```
# This is a comments
```

Variables:

Variables are generally declared using the = symbol and no spaces before and after it. Variables are accessed using the \$ symbol, or (better) using also curly brackets {}

Example: `set_a_var.sh`

```
#!/bin/bash
my_var="Ciao!"
my_list="Luke Leia Anakin Obi-Wan"
echo "The value of my_var is: ${my_var}"
echo "The value of my_list is: ${my_list}"
```

```
giorgio@Kubuntu-22:~/scripts$ ./set_vars.sh
The value of my_var is: Ciao!
The value of my_list is: Luke Leia Anakin Obi-Wan
giorgio@Kubuntu-22:~/scripts$ █
```

Variables: "Arithmetic expansion"

Bash allows for some maths on integers. For floating-point numbers, you can pipe the expression to **bc** (basic calculator).

`$((expression))`: This operator called "arithmetic expansion". It is used to perform some (integer-based) maths in Bash. `(())` evaluates the expression, `$` stores the result

Example:

```
a=2
echo "$((${a}**2))" # prints a*a = 4
```

```
b=2.5
echo "${b}^2" | bc -l # prints b*b = 6.25
echo "$((${b}**2))" # will cause an error: no integer!
```

```
i=10
echo "$((++i))" # will return 11
```

| Syntax | Description |
|---------------------------------------|---|
| <code>++x, x++</code> | Pre and post-increment |
| <code>--x, x--</code> | Pre and post-decrement |
| <code>+, -, *, /</code> | Addition, subtraction, multiplication, division |
| <code>%, ** or ^</code> | Modulo (remainder) and exponentiation |
| <code>&&, , !</code> | Logical AND, OR, and negation |
| <code>&, , ^, ~</code> | Bitwise AND, OR, XOR, and negation |
| <code><=, <, >, >=</code> | Less than or equal to, less than, greater than, and greater than or equal to comparison operators |
| <code>==, !=</code> | Equality and inequality comparison operators |

User input: User input (e.g. from the keyboard) can be assigned to a variable using **read**.

Example: **origin.sh**

```
#!/bin/bash
echo "From which country do you come from?"
read country
echo "You come from ${country}"
```

```
giorgio@Kubuntu-22:~/scripts$ ./origin.sh
From which country do you come from?
SOKOVIA
You come from SOKOVIA
giorgio@Kubuntu-22:~/scripts$
```

Bash arguments: You can pass arguments to your Bash scripts. They can be accessed from the script using **\$1, \$2, \$3, ..., \$n**, with n the order they are passed. **\$@** is a reference to ALL passed arguments.

Example: **fruit_salad.sh**

```
#!/bin/bash
echo "First fruit is $1"
echo "Second fruit is $2"
echo "Third fruit is $3"
echo "All fruits are $@"
```

```
giorgio@Kubuntu-22:~/scripts$ ./fruit_salad.sh apples kiwis bananas
First fruit is apples
Second fruit is kiwis
Third fruit is bananas
All fruits are apples kiwis bananas
giorgio@Kubuntu-22:~/scripts$
```

Arrays: An array is initialised by assigning values separated by spaces and enclosed in round parentheses **()**. You can access array values in different ways.

Example: **arrays.sh**

```
my_array=("X-Wing" "A-Wing" "B-Wing" "Y-Wing")
echo "All spaceships: ${my_array[@]}"
echo "Indices of items are: ${!my_array[@]}"
echo "Number of items in the array is: $#my_array[@]"
echo "First spaceship is: ${my_array[0]}" # 0-index based!
echo "Second spaceship is: ${my_array[1]}"
echo "Last spaceship is: ${my_array[-1]}"
echo "First two spaceships are: ${my_array[@]:0:2}" # 2 is excluded
```

```
giorgio@Kubuntu-22:~/scripts$ ./arrays.sh
All spaceships: X-Wing A-Wing B-Wing Y-Wing
Indices of items are: 0 1 2 3
Number of items in the array is: 4
First spaceship is: X-Wing
Second spaceship is: A-Wing
Last spaceship is: Y-Wing
First two spaceships are: X-Wing A-Wing
giorgio@Kubuntu-22:~/scripts$
```

Conditional expressions: the `[[` compound command and the `[]` built-in command are used to test file attributes and perform string and arithmetic comparisons.

(Some) examples of file expressions:

`[[-e $file]]` : returns true if file exists

`[[-d $file]]` : returns true if file exists and is a directory

`[[-x $file]]` : returns true if file is executable

(Some) examples of string expressions:

`[[$string1 == $string2]]` : returns true if the strings are equal

`[[$string1 != $string2]]` : returns true if the strings are different

Conditional expressions (ctd)

(Some) examples of arithmetic operators:

`[[${arg1} -eq ${arg2}]]` : returns true if the 2 numbers are equal

`[[${arg1} -ne ${arg2}]]` : returns true if the 2 numbers are different

`[[${arg1} -gt ${arg2}]]` : returns true if arg1 is greater than arg2

`[[${arg1} -le ${arg2}]]` : returns true if arg1 is less or equal than arg2

`[[test_case_1]] && [[test_case_2]]` : returns true if both cases are true (AND)

`[[test_case_1]] || [[test_case_2]]` : returns true if at least one of the cases is true (OR)

Examples of Exit status operators:

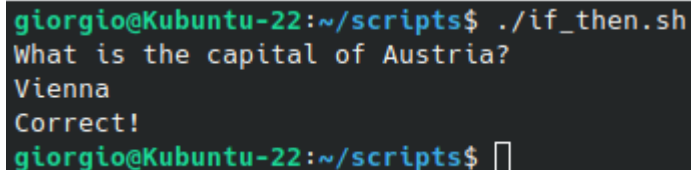
`[[$? -eq 0]]` : returns true if the command was successful without any errors

`[[$? -gt 0]]` : returns true if the command was not successful or had errors

Conditional statements: The conditional expressions seen in the previous slides can be used to build conditional statements such as *if-then*, *if-then-else*, etc.

Example: **if_then.sh**

```
#!/bin/bash
correct_answer="Vienna"
echo "What is the capital of Austria?"
read answer
if [ ${answer} == ${correct_answer} ]
then
    echo "Correct!"
fi # please note the "fi" to close the statement
```



```
giorgio@Kubuntu-22:~/scripts$ ./if_then.sh
What is the capital of Austria?
Vienna
Correct!
giorgio@Kubuntu-22:~/scripts$
```

Conditional statement (ctd)

Example: `if_then_else.sh`

```
#!/bin/bash
correct_answer="Glera"
echo "Prosecco wine is made with grapes named...?"
read answer
if [ ${answer} == ${correct_answer} ]
then
    echo "Correct! You deserve a glass of Prosecco! :-)"
else
    echo "Oh no, you seem to lack some basic knowledge :-("
fi
```

```
giorgio@Kubuntu-22:~/scripts$ ./if_then_else.sh
Prosecco wine is made with grapes named...?
Glera
Correct! You deserve a glass of Prosecco! :-)
giorgio@Kubuntu-22:~/scripts$ ./if_then_else.sh
Prosecco wine is made with grapes named...?
Sauvignon
Oh no, you seem to lack some basic knowledge :-)
```

Bash shell scripting

Conditional statement (ctd)

Example: **case.sh**

```
#!/bin/bash
echo "Enter the name of a Pixar movie:"
read -r pixar_movie # -r allows to have spaces in input
case ${pixar_movie} in
    "Toy Story 5")
        echo "${pixar_movie} will be released in 2026"
        ;;
    "Toy Story" | "Toy Story 2" | "Toy Story 3" | "Toy Story 4")
        echo "${pixar_movie} tells the adventures of Woody, Buzz etc."
        ;;
    "Monsters Inc.")
        echo "With ${pixar_movie} you'll fall in love with Boo!"
        ;;
    *)
        echo "At the moment, I can't tell you anything about ${pixar_movie}..."
        ;;
esac
```

```
giorgio@Kubuntu-22:~/scripts$ ./case.sh
Enter the name of a Pixar movie:
Toy Story 5
Toy Story 5 will be released in 2026
giorgio@Kubuntu-22:~/scripts$ ./case.sh
Enter the name of a Pixar movie:
Toy Story 3
Toy Story 3 tells the adventures of Woody, Buzz etc.
giorgio@Kubuntu-22:~/scripts$ ./case.sh
Enter the name of a Pixar movie:
Monsters Inc.
With Monsters Inc. you'll fall in love with Boo!
giorgio@Kubuntu-22:~/scripts$ ./case.sh
Enter the name of a Pixar movie:
Wall-E
At the moment, I can't tell you anything about Wall-E...
giorgio@Kubuntu-22:~/scripts$ █
```

Linux in a nutshell
Terminal & shell(s)
File system
File system
permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
**Bash shell
scripting**
Further resources

Bash shell scripting

Loop statements: In Bash there are loops, while-loops, and until-loops. Beware: you loop over lists! Arrays, if applicable, must be cast to lists.

Example: **loop.sh**

```
#!/bin/bash
fruits_list="apples kiwis bananas strawberries"
fruits_array=("mangos" "peaches" "apricots" "pears")
echo "*** Iterating over a list"
for fruit in ${fruits_list}
do
    echo "Printing: ${fruit}"
done
echo -e "\n"
echo "*** Iterating over an array"
for fruit in ${fruits_array[@]} # Array cast to list!
do
    echo "Printing: ${fruit}"
done
echo -e "\n"
```

```
giorgio@Kubuntu-22:~/scripts$ ./loop.sh
*** Iterating over a list
Printing: apples
Printing: kiwis
Printing: bananas
Printing: strawberries

*** Iterating over an array
Printing: mangos
Printing: peaches
Printing: apricots
Printing: pears

giorgio@Kubuntu-22:~/scripts$
```

Linux in a nutshell
Terminal & shell(s)
File system
File system
permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
**Bash shell
scripting**
Further resources

Example: **while_until_loop.sh**

```
#!/bin/bash
echo "*** Example of a while-loop"
counter=1
while [[ $counter -le 5 ]]
do
    echo "While-loop counter is: ${counter}"
    ((counter++))
done

echo -e "\n*** Example of an until-loop"
counter2=1
until [[ $counter2 -gt 5 ]]
do
    echo "Until-loop counter is: ${counter2}"
    ((counter2++))
done
```

```
giorgio@Kubuntu-22:~/scripts$ ./while_until.sh
*** Example of a while-loop
While-loop counter is: 1
While-loop counter is: 2
While-loop counter is: 3
While-loop counter is: 4
While-loop counter is: 5

*** Example of an until-loop
Until-loop counter is: 1
Until-loop counter is: 2
Until-loop counter is: 3
Until-loop counter is: 4
Until-loop counter is: 5
giorgio@Kubuntu-22:~/scripts$
```

Loop statements (ctd):

Inside a loop statement, you can use commands `continue` and `break`. With **continue** you can stop the current iteration of the loop and start with the next one.

Example: `continue.sh`

```
#!/bin/bash
for n in 1 2 3 4 5
do
  if [ ${n} == 3 ]
  then
    echo "Skipping value ${n}"
    continue
  else
    n_squared=$((n*n)) # $(( )) is called "arithmetic expansion"
    echo "Current value is: ${n}, its square is: ${n_squared}"
  fi
done
```

```
giorgio@Kubuntu-22:~/scripts$ ./break_continue.sh
Current value is: 1, its square is: 1
Current value is: 2, its square is: 4
Skipping value 3
Current value is: 4, its square is: 16
Current value is: 5, its square is: 25
giorgio@Kubuntu-22:~/scripts$
```

Loop statements (ctd):

With **break** you can exit a loop if a certain condition is met.

Example: **break.sh**

```
#!/bin/bash
for n in 1 2 3 4 5
do
  if [ ${n} == 3 ]
  then
    echo "Skipping value ${n}"
    echo -e "\nExiting loop!"
    break
  else
    n_squared=$((n*$n)) # $(( )) is called "arithmetic expansion"
    echo "Current value is: ${n}, its square is: ${n_squared}"
  fi
done
```

```
giorgio@Kubuntu-22:~/scripts$ ./break.sh
Current value is: 1, its square is: 1
Current value is: 2, its square is: 4
Skipping value 3

Exiting loop!
giorgio@Kubuntu-22:~/scripts$
```

Index

- Linux in a nutshell
- The terminal and the shell(s)
- File system
- File system permissions
- Data streams
- Processes
- Scheduling jobs
- Installation of software applications
- Text editors
- Bash shell scripting
- **Further resources**

Further resources

Linux in a nutshell
Terminal & shell(s)
File system
File system
permissions
Data streams
Processes
Scheduling jobs
Software install
Text editors
Bash shell scripting
Further resources

- This Introduction is part of **TU Delft's GeoGeeks**
 - <https://tudelft3d.github.io/geogeeks/>
- **"Introduction to Linux"**, by the Linux Foundation
 - <https://training.linuxfoundation.org/training/introduction-to-linux/>
- **"Linux Tutorial"**, by GeeksforGeeks
 - <https://www.geeksforgeeks.org/linux-tutorial/>
- **"Introduction to Linux"**, by CodeAcademy
 - <https://www.codecademy.com/learn/introduction-to-linux>



Dr. Giorgio Agugiaro

g.agugiaro@tudelft.nl

3D Geoinformation Group

TU Delft

The Netherlands

<https://3d.bk.tudelft.nl/gagugiaro>

Acknowledgements

Some parts of this document take inspiration and expand previous material by Clara García-Sánchez and Akshay Patil (TU Delft), originally developed for [GeoGeeks](#).