

Linux A Visual Guide

Essential Concepts and Practical Skills for New Systems Engineers

Reiner Braun

Sample Book — Chapters 1 to 2

Introduction

This book is designed for readers who are taking their first steps with Linux, including systems engineers who want to build the right foundation from the start.

Rather than asking you to memorize more than you need, the book focuses on the essential concepts that will support your future study and day-to-day work. More advanced topics are included only when they are likely to become useful to many readers, and even then they are introduced only far enough to explain their purpose, context, and overall role.

Processes, relationships, and structures that can be difficult to picture from text alone are supported by diagrams. By moving between explanations, visuals, and examples, you can build a clearer and more practical understanding of the subject.

Product names, service names, logos, icons, and other identifying marks mentioned in this book are trademarks or registered trademarks of their respective owners. This book is not an official publication, certified training resource, or endorsed learning material of any company or organization.

How to Use This Book

Read the chapters in order. The sequence builds the necessary background step by step. Later chapters assume that you understand the command line, files, identities, permissions, processes, packages, services, and other foundations introduced earlier.

Preview what the chapter will enable you to do. At the beginning of each chapter, the opening explanation serves as **What You'll Be Able to Do**: it states the chapter's central conclusion, operating model, and practical outcome. Review that opening before working through the details so you know what capability you are building.

Use diagrams to understand relationships, processes, and structures. A diagram is a companion to the surrounding explanation, not a replacement for it. Follow the arrows, boundaries, containment, and sequence labels, then return to the text to connect the visual model to Linux commands and evidence.

Try the commands and hands-on examples whenever possible. Use a non-production Linux virtual machine or another disposable lab environment. Read the stated prerequisites, predict what a command will do, run the smallest example, compare the actual result with the expected result, and restore the lab when necessary.

Confirm the fundamentals before moving on. Use each chapter's concluding summary to restate the main idea in your own words, then answer these review questions:

- What system layer, resource, or responsibility does the chapter explain?
- What evidence should you inspect before making a change?
- What common beginner mistake could lead to an incorrect conclusion or unsafe action?
- How would you verify that an action produced the intended result?

Treat advanced-topic overviews as signposts. You do not need to memorize their commands or internal details at this stage. Remember what problem each topic solves, where it fits in the Linux system, and when you may need to study it more deeply.

Table of Contents

[Introduction](#)

[How to Use This Book](#)

[The Foundation in One View](#)

[Chapter 1: Linux as a Layered System](#)

[Start with the Layers](#)

[What the Kernel Actually Does](#)

[User Space, Services, and Applications](#)

[Terminal, Shell, and Command-Line Program](#)

[Why Distributions Differ](#)

[Identify an Unfamiliar Host](#)

[Runtime Interfaces: /dev, /proc, and /sys](#)

[A Working Mental Model](#)

[Chapter 2: The Shell and the Anatomy of a Command](#)

[The Shell Is an Interpreter, Not the Operating System](#)

[Reading the Prompt Carefully](#)

[Commands, Options, and Arguments](#)

[Case and Whitespace Matter](#)

[Quoting and Escaping](#)

[Wildcards Are Expanded by the Shell](#)

[What a Command Name Can Mean](#)

[How PATH Controls Command Lookup](#)

[Variables and the Process Environment](#)

[Interactive, Login, and Non-Interactive Shells](#)

[Exit Status Is the Primary Result Signal](#)

[Discover Commands Safely](#)

[Completion and History as Accuracy Tools](#)

[There Is No Universal Undo](#)

[A Practical Command-Discovery Sequence](#)

The Foundation in One View

A systems engineer does not become effective by memorizing a long list of Linux commands. The essential skill is understanding what layer of the system is involved, inspecting the current state, making the smallest justified change, and verifying the result.

Linux systems are built from a few recurring ideas. The kernel controls hardware and system resources. Programs run as processes with particular identities and permissions. Files and directories form one tree, even when the data comes from different devices. Services are long-running processes managed by the operating system. Administrative tools expose these layers through text-based interfaces that can be inspected and combined.

This guide begins by establishing that operating model through the command line, files, text processing, identities, permissions, processes, packages, boot, and services. These subjects are tightly connected. For example, a service may fail because its package is missing a library, its process lacks permission to read a configuration file, or its systemd unit points to the wrong working directory. A correct mental model lets you distinguish these causes instead of trying unrelated fixes.

Practice should take place on a **non-production Linux virtual machine** where mistakes can be reversed. Some examples require administrative privileges through `sudo`; those examples state the requirement before the command. Commands that remove data, replace files, or change system configuration deserve a deliberate pause before execution.

Part I: Build the Linux Mental Model

Chapter 1: Linux as a Layered System

Start with the Layers

Linux is easiest to understand as a set of layers that cooperate through controlled interfaces. From the bottom upward, a typical system contains hardware, the Linux kernel, user-space programs and libraries, background services, applications, and user interfaces.

The word **Linux** is used in two related ways. Strictly, the **Linux kernel** is the core program that manages hardware and system resources. In everyday use, “Linux” often means a complete operating system built around that kernel. A **Linux distribution**, or distro, is a packaged operating system that combines the Linux kernel with command-line tools, libraries, an installer, a package manager, default configuration, and a software repository.


Ubuntu, Debian, Red Hat Enterprise Linux, Rocky Linux, AlmaLinux, Fedora, SUSE Linux Enterprise Server, and openSUSE are examples of distributions. They share many concepts and commands, but they may package software differently, store some configuration in different locations, or select different administrative tools.

A useful layered view is:

1. **Hardware** provides processors, memory, storage devices, network adapters, and other physical or virtual components.
2. **The kernel** controls access to that hardware and allocates resources.
3. **User space** contains ordinary programs, system libraries, and services that run outside the kernel.
4. **Services and applications** perform useful work such as accepting SSH connections, hosting a website, or running a database.
5. **User interfaces** let people and automation interact with the system through a shell, a graphical interface, or a remote protocol.

Chapter 1: Linux as a Layered System

1-1 Linux System Layers and Responsibility Boundaries



Applications request; the kernel controls access.

001

Figure 1-1: Linux System Layers and Responsibility Boundaries


Think of a large office building. The kernel is the building manager: it assigns rooms, controls elevators, tracks access to shared facilities, and responds to maintenance requests. Applications are tenants. They do not rewire the building directly; they submit controlled requests through approved channels.

What the Kernel Actually Does

The kernel is responsible for several fundamental jobs:

- **CPU scheduling:** deciding which runnable process uses a processor and for how long.
- **Memory management:** assigning virtual memory to processes and protecting one process from another.
- **Device management:** communicating with disks, network adapters, terminals, and other devices through drivers.
- **Filesystem management:** presenting stored data through files and directories.
- **Process management:** creating, tracking, and terminating running programs.
- **Networking:** implementing protocols, interfaces, routes, and sockets.
- **Security enforcement:** applying identities, permissions, and other policy checks.

Chapter 1: Linux as a Layered System 1-2 Kernel Responsibility Map



002

Figure 1-2: Kernel Responsibility Map

A program normally cannot touch hardware or kernel memory directly. Instead, it uses a **system call**, which is a controlled request from a user-space program to the kernel. Opening a file, creating a process, allocating memory, and sending network data all ultimately involve system calls.

For example, when the `cat` program reads a file, the sequence is conceptually:


1. The shell starts `cat` as a process.
2. `cat` asks the kernel to open the named file.
3. The kernel checks the process identity and file permissions.
4. If access is allowed, the kernel reads data from the filesystem and returns it to `cat`.

5. `cat` asks the kernel to write that data to its standard output, usually the terminal.

Chapter 1: Linux as a Layered System

1-2 What Happens When `cat` Reads a File

Applications request; the kernel decides and mediates.



003

Figure 1-3: What Happens When `cat` Reads a File

This separation is important. **Applications request resources; the kernel decides whether and how those requests are fulfilled.**

User Space, Services, and Applications


User space is the part of the operating environment in which ordinary programs run. The name does not mean that every program is started directly by a human. It includes interactive commands, programming-language runtimes, system libraries, web servers, databases, monitoring agents, and many background services.

A **process** is a running instance of a program. The executable file on disk is passive; it becomes a process when the kernel loads it and begins executing it. Several processes can run from the same executable at the same time.

A **service** is a long-running function provided by one or more processes, usually started and supervised by the operating system. An SSH service accepts remote login connections. A web service listens for HTTP requests. A logging service collects or stores event messages. The older term **daemon** is commonly used for a background service process.

Chapter 1: Linux as a Layered System

1-3 Executable, Process, and Service Relationship



004

Figure 1-4: Executable, Process, and Service Relationship

Linux is a **multiuser** system: it can maintain separate identities, permissions, and sessions for multiple users. It is also a **multitasking** system: it rapidly schedules many processes so they can make progress at the same time. Even a small server may run hundreds of processes, most of which spend much of their time waiting for input, timers, storage, or network traffic.

Terminal, Shell, and Command-Line Program

These terms are related but not interchangeable.

A **terminal** is the text-oriented interface that displays output and sends keyboard input. On a desktop, a terminal emulator is a graphical application that behaves like a traditional hardware terminal. An SSH session also provides a remote terminal.

A **shell** is a program that reads command text, interprets shell syntax, and starts other programs. Bash is common, but other shells include Dash, Zsh, Fish, and Ksh.

A **command-line program** is an executable tool such as `ls`, `grep`, or `systemctl`. The shell may interpret part of the command line first, but the external program performs its own operation after it starts.

A **graphical interface** is another user-space application. It does not replace the kernel or create a separate operating system; it uses the same underlying processes, files, devices, and permissions.


Consider this command:

```
ls -l /var/log
```

The terminal sends the text to the shell. The shell identifies `ls` as the command and starts it. The `ls` process asks the kernel for directory information. The kernel checks access, returns metadata, and the process writes formatted output back to the terminal.

Chapter 1: Linux as a Layered System

1-4 Terminal, Shell, Command, Kernel, and Output Flow



005

Figure 1-5: Terminal, Shell, Command, Kernel, and Output Flow

Why Distributions Differ

Distributions select and integrate the software around the kernel. Differences commonly include:

- The package format and package manager.
- Default versions of applications and libraries.
- Configuration defaults and enabled services.
- Security frameworks and firewall frontends.
- Support lifecycle and update policy.
- Locations used by a particular packaged application.

These differences explain why a command copied from one server may not work unchanged on another. For example, Debian and Ubuntu commonly use `apt` for package management, while Red Hat-compatible systems commonly use `dnf`. The underlying task is the same, but the interface differs.

Do not begin administrative work by assuming the distribution, release, architecture, or shell. Identify them first.

Identify an Unfamiliar Host

The following inspection requires only an ordinary shell account and makes no system changes:

```
whoami
uname -sr
uname -m
cat /etc/os-release
printf 'Login shell: %s\n' "$SHELL"
ps -p $$ -o comm=
```

The commands answer different questions:

- `whoami` prints the current username.
- `uname -sr` prints the kernel name and kernel release.
- `uname -m` prints the machine architecture, such as `x86_64` or `aarch64`.
- `/etc/os-release` contains distribution identification and release information.
- `$SHELL` usually names the account's configured login shell.
- `ps -p $$ -o comm=` prints the name of the shell process currently running this command.

Expected results vary by host. A machine might report Ubuntu 24.04 with a Linux 6.x kernel on `x86_64`, while another reports Rocky Linux on `aarch64`. **The distribution release and the kernel release are separate facts.** A distribution can update its kernel without changing its overall release name.

The `$SHELL` value and the current shell can also differ. For example, an account may have Bash configured as its login shell but currently be running a nested `sh` session. That distinction matters because shell syntax and startup behavior can differ.

The `arch` command is another common way to display machine architecture:

```
arch
```

It usually reports the same architecture name as `uname -m`.

Runtime Interfaces: `/dev`, `/proc`, and `/sys`


Some paths in the Linux directory tree do not behave like ordinary stored files.

`/dev` contains **device nodes**, which are special filesystem entries through which programs access devices or kernel-provided interfaces. For example, `/dev/null` discards data written to it, and terminal devices appear under `/dev/pts` on many systems.

`/proc` is a virtual filesystem that exposes process and kernel information. It is generated at runtime rather than stored as ordinary disk files. Each process normally has a directory named after its process ID, such as `/proc/1234`.

Chapter 1: Linux as a Layered System

1-7 Runtime Interfaces: /dev, /proc, and /sys



006

Figure 1-6: Runtime Interfaces: /dev, /proc, and /sys

On a typical Linux system, an ordinary user can inspect a few safe examples:

```
cat /proc/version
cat /proc/meminfo | head
ls /sys/class/net
```

These commands display kernel version details, the first lines of memory information, and the names of network interfaces known to the kernel. The exact output is host-specific. Many files under /proc and /sys are read-only to ordinary users; some writable entries control live kernel behavior and should not be changed casually.

Treat /proc and /sys as live views and controls, not as ordinary configuration storage. Their contents are recreated as the system runs, and persistent configuration is generally managed elsewhere.

A Working Mental Model


When you encounter an unfamiliar Linux problem, first ask which layer owns the behavior:

- Is the kernel reporting a device or resource problem?
- Is a user-space process missing, blocked, or failing?
- Is a service manager starting the process with unexpected settings?
- Is a distribution-specific package or configuration involved?
- Is the shell interpreting the command differently than expected?

- Is an identity or permission preventing access?

Chapter 1: Linux as a Layered System

1-8 Map a Linux Symptom to the Responsible Layer



007

Figure 1-7: Map a Linux Symptom to the Responsible Layer

This layered question does not solve every problem immediately, but it prevents random changes. **Linux administration becomes manageable when each symptom is mapped to the smallest relevant layer.**

Chapter 2: The Shell and the Anatomy of a Command

The Shell Is an Interpreter, Not the Operating System


A shell turns command text into actions by applying its own syntax before it starts programs. Understanding that order is more valuable than memorizing isolated commands because many command-line mistakes happen before the intended program even begins.

A shell normally performs work in stages:

1. It reads a line of text.
2. It recognizes shell syntax such as quotes, variables, wildcards, pipelines, and redirections.
3. It expands or transforms parts of the line.
4. It determines what command name refers to.
5. It starts the command or runs a shell built-in.
6. It waits for or tracks the result and records an exit status.

Chapter 2: The Shell and the Anatomy of a Command

2-1 Shell Interpretation Stages



Typed text is not necessarily what the program receives.

008

Figure 2-1: Shell Interpretation Stages

This is why the same program can receive different arguments depending on how the command line is quoted. It is also why a redirection failure can occur before the program starts.

Reading the Prompt Carefully

An interactive shell usually displays a **prompt**, which is text indicating that it is ready for a command. Prompt formats vary. A common convention is:

```
alice@server:~$
```

This example suggests the user is `alice`, the host is `server`, the current directory is the home directory represented by `~`, and the account is not root. A root prompt often ends with `#` rather than `$`:

```
root@server:~#
```

These characters are conventions, not security controls. A prompt can be customized or misleading. Verify identity with `whoami` or `id` when authority matters.

In books and documentation, `$` and `#` are often shown as prompt markers. They are normally not part of the command. For example:

```
$ id
```

means to type `id`, not `$ id`.

Before running a copied command, confirm the current user, current directory, and target path. Those three facts determine the effect of many administrative commands.

Commands, Options, and Arguments

Consider:

```
ls -l --human-readable /var/log
```

The command line contains:

- `ls`: the command name.
- `-l`: a short option requesting a long listing.
- `--human-readable`: a long option requesting readable size units where applicable.
- `/var/log`: a path argument identifying the directory to list.

Chapter 2: The Shell and the Anatomy of a Command

2-3 Command, Option, and Argument Anatomy




Figure 2-2: Command, Option, and Argument Anatomy

An **option** modifies a command's behavior. An **argument** supplies data or identifies a target. Some options require their own value, as in:

```
grep --max-count=3 'failed' system.log
```

Here, `--max-count=3` limits matches, `'failed'` is the search pattern, and `system.log` is the input filename.

Command syntax is defined by each command, not by one universal Linux rule. Some commands accept combined short options, such as `ls -la`; some do not. Some accept options after operands; others expect options first. Read that command's help rather than guessing.

The marker `--` commonly means **end of options**. It is useful when a filename begins with a hyphen:

```
rm -- -old-report
```

Prerequisite: a file literally named `-old-report` in the current directory. The command tells `rm` to treat the name as a file operand rather than an option. The expected result is removal of that file after normal permission checks. Use this only after confirming the name with a safe listing such as `ls -l -- -old-report`.

Case and Whitespace Matter

Linux filenames and command names are normally case-sensitive. `Report.txt`, `report.txt`, and `REPORT.TXT` can be three different files. Environment-variable names are also case-sensitive: `$PATH` and `$Path` are unrelated.

Whitespace separates words after the shell has processed quoting and escaping. These commands are not equivalent:

```
printf '%s\n' one two
printf '%s\n' 'one two'
```

The first passes two data arguments, `one` and `two`. The second passes one argument containing a space. Both commands use `printf` to print each argument on its own line.

Repeated unquoted spaces usually act as one separator. Quoted spaces are preserved as data.

Quoting and Escaping

Quoting controls what the shell interprets.

Single quotes preserve literal text

Inside single quotes, almost every character is treated literally:

```
printf '%s\n' '$HOME * ?'
```

Expected output:

```
$HOME * ?
```

The variable and wildcard characters are not expanded.

Double quotes allow selected expansion

Inside double quotes, variable expansion and command substitution still occur, but spaces in the expanded result remain part of one argument:

```
name='Project Alpha'
printf 'Name: %s\n' "$name"
```

Expected output:

```
Name: Project Alpha
```

Quoting `"$name"` ensures that the value is passed as one argument. **As a general shell habit, quote variable expansions unless you specifically require word splitting or wildcard expansion.**

A backslash escapes one character

A backslash can remove the special meaning of the next character:

```
printf '%s\n' Project\ Alpha
```

This passes `Project Alpha` as one argument. Quoting is usually clearer when a value contains several special characters.

Single quotes cannot directly contain a single quote. A common shell form for text such as `don't` is:


```
printf '%s\n' 'don''''t'
```

This combines three quoted pieces into one argument. Such syntax is valid but difficult to read; for ordinary fixed text, double quotes may be simpler when no unwanted expansion is present.

Chapter 2: The Shell and the Anatomy of a Command

2-5 How Quoting and Escaping Change Shell Expansion

Assume `HOME=/home/alice`; files: `app.log`, `error.log`



010

Figure 2-3: How Quoting and Escaping Change Shell Expansion

Wildcards Are Expanded by the Shell

A **glob** is a shell pattern used to match filenames. Common wildcard characters are:

Pattern	Meaning in ordinary shell globbing
*	Zero or more characters
?	Exactly one character
[abc]	One character from the listed set
[0-9]	One character from the stated range

Suppose the current directory contains `app.log`, `auth.log`, and `notes.txt`. Then:


```
printf '%s\n' *.log
```

is normally transformed by the shell into something like:

```
printf '%s\n' app.log auth.log
```

Chapter 2: The Shell and the Anatomy of a Command

2-6 Shell Wildcard Expansion



011

Figure 2-4: Shell Wildcard Expansion

The `printf` program does not receive the pattern; it receives the matched filenames. Quoting the pattern prevents expansion:

```
printf '%s\n' '*.log'
```

Expected output is the literal text `*.log`.

A pattern that matches nothing is handled differently by different shells or shell settings. Bash commonly leaves an unmatched pattern unchanged. Do not assume that a wildcard always produces at least one filename.

Hidden names beginning with `.` are not normally matched by a leading `*`. This is one reason `rm -rf *` is both dangerous and incomplete: it can delete most visible entries while leaving hidden ones.

What a Command Name Can Mean


A typed command name may refer to several kinds of shell object:

- A **shell built-in** is implemented inside the shell itself, such as Bash's `cd`.
- An **executable file** is an external program stored in the filesystem.
- An **alias** is a shell-defined text replacement, often used for convenience.

- A **shell function** is a named block of shell code in the current shell environment.

Chapter 2: The Shell and the Anatomy of a Command

2-7 What a Command Name Can Resolve To



012

Figure 2-5: What a Command Name Can Resolve To

The distinction matters. `cd` must change the working directory of the current shell, so it cannot work as an ordinary child process that exits immediately.

Use `type` in an interactive Bash-like shell to learn how a name resolves:

```
type cd
type ls
type -a ls
```

The first command should report that `cd` is a shell built-in. The result for `ls` may show an alias, an executable path, or both. `type -a` lists all known resolutions.

`command -v` provides a script-friendly way to check whether and how a command is available:

```
command -v systemctl
```

Expected output on a system with `systemd` is typically a path such as `/usr/bin/systemctl`. A nonzero exit status means the shell could not resolve the name according to its command-search rules.

How `PATH` Controls Command Lookup


`PATH` is an environment variable containing a colon-separated list of directories that the shell searches for executable commands:

```
printf '%s\n' "$PATH"
```

A typical value might include `/usr/local/bin`, `/usr/bin`, and `/bin`. When you type `grep`, the shell searches these directories in order and runs the first matching executable it finds, unless a built-in, function, or alias takes precedence.

Chapter 2: The Shell and the Anatomy of a Command

2-8 How PATH Selects an Executable



013

Figure 2-6: How PATH Selects an Executable

To display each path entry on a separate line without changing anything:

```
printf '%s\n' "$PATH" | tr ':' '\n'
```

Expected output is one search directory per line.

The current directory is not always in `PATH`, and excluding it is a useful safety measure. To run an executable file named `check.sh` from the current directory, use an explicit path:

```
./check.sh
```

The `./` means "the file named `check.sh` in the current directory." This avoids accidentally running an unrelated program found elsewhere in `PATH`.

Do not add untrusted or world-writable directories near the front of an administrator's `PATH`. A malicious executable with a familiar name could be selected before the intended system command.

Variables and the Process Environment

A **shell variable** is a name and value maintained by the current shell. An **environment variable** is a name and value exported so that newly started child processes receive a copy.

Create and inspect a shell variable:

```
project='inventory'
printf '%s\n' "$project"
```

The assignment contains no spaces around `=`. Expected output is `inventory`.

A child shell does not automatically receive an unexported variable:

```
project='inventory'
sh -c 'printf "child sees: <%s>\n" "$project"'
```


The child will normally print an empty value. Exporting changes that:

```
export project
sh -c 'printf "child sees: <%s>\n" "$project"'
```

Expected output now includes inventory.

Chapter 2: The Shell and the Anatomy of a Command

2-9 Shell Variables and Process Environment Inheritance



014

Figure 2-7: Shell Variables and Process Environment Inheritance

The environment is copied when a process starts. Later changes in the parent do not rewrite the environment of an already-running child. Similarly, a child process cannot directly change the parent shell's variables or current directory.

Use `env` to display the current process environment, but remember that environment values can contain secrets. Avoid posting complete environment output into tickets or chat without review.

Interactive, Login, and Non-Interactive Shells

A shell can run in different modes:


- An **interactive shell** reads commands from a user and displays a prompt.
- A **non-interactive shell** reads commands from a script or a supplied command string.
- A **login shell** represents the start of a login session and may read a different set of startup files.

The exact startup files depend on the shell and how it is invoked. In Bash, login shells and non-login interactive shells commonly read different configuration files. Non-interactive scripts may not read the files that create aliases or modify `PATH` in your terminal.

This explains a common operational problem: a command works when typed manually but fails in a script, scheduler, or service. The automated context may have a different working directory, `PATH`, environment, user, or shell.

Chapter 2: The Shell and the Anatomy of a Command

2-10 Shell Modes and Differing Execution Contexts



015

Figure 2-8: Shell Modes and Differing Execution Contexts

A shell environment is context, not a universal property of the host. When diagnosing a discrepancy, compare the actual user, shell, environment, and invocation method.

Exit Status Is the Primary Result Signal

When a command finishes, it returns an **exit status**, which is a small integer reported to its parent process. By convention, `0` usually means success and a nonzero value indicates some form of failure or exceptional result. The exact nonzero meanings are command-specific.

In an interactive shell, `$?` expands to the most recent command's exit status:

```
true
printf 'status=%s\n' "$?"
false
printf 'status=%s\n' "$?"
```

Expected output is commonly:

```
status=0
status=1
```


Check `$?` immediately, because every subsequent command replaces it. In scripts, a clearer form is often an `if` statement:

```
if grep -q 'root' /etc/passwd; then
    printf '%s\n' 'entry found'
else
    printf '%s\n' 'entry not found or read failed'
fi
```

This example uses `grep -q` to search without printing matches. The shell chooses a branch based on the exit status.

Chapter 2: The Shell and the Anatomy of a Command

2-11 Exit Status and Conditional Control Flow



016

Figure 2-9: Exit Status and Conditional Control Flow

Visible output is not a reliable success indicator. Some successful commands print nothing. Some failed commands produce output on standard output, and some informational commands intentionally return nonzero when a condition is false rather than because the program malfunctioned.

Discover Commands Safely

Linux systems include local documentation. Use it before relying on an unfamiliar example.

Concise built-in help

Many programs support `--help`:

```
ls --help
```

This usually prints a compact option summary. For a shell built-in, Bash provides `help`:

```
help cd
```

Manual pages

The `man` command opens a **manual page**, which is structured local documentation:

```
man ls
```

Inside the common `less` viewer, press `/` to search, `n` for the next match, and `q` to quit.

Manual pages are organized into numbered sections. The same name can exist in more than one section. Frequently encountered sections include:

Section	Typical content
1	User commands
2	System calls
3	Library functions
5	File formats and configuration files
8	System administration commands

For example, `man 5 passwd` describes the `/etc/passwd` file format, while `man 1 passwd` describes the command used to change passwords.

Use `man -k` to search manual-page descriptions:

```
man -k 'copy files'
```

The expected result is a list of potentially relevant manual pages when the manual-page index is installed and current.

Info documentation

Some GNU tools provide detailed `info` pages:

```
info coreutils
```

The interface is more extensive than a simple manual page. It is useful when `man` gives only a summary.

Completion and History as Accuracy Tools

Tab completion asks the shell to complete command names, filenames, or options where supported. It reduces typing errors and exposes available paths. Pressing Tab twice often shows multiple possible completions.

Command history lets the shell recall earlier commands. In Bash, the Up and Down arrow keys navigate recent entries, and `Ctrl-r` starts a reverse history search. History is useful, but an old command may contain obsolete paths, dangerous options, or credentials. Review the full line before running it again.

Avoid placing secrets directly on command lines. They may be visible in shell history, process listings, audit records, or terminal logs. Prefer a tool's protected prompt, a restricted credential file, or another mechanism designed for secrets.

There Is No Universal Undo

Graphical applications often provide an Undo command. The Linux command line does not have a universal equivalent. A command may alter a file, remove a directory, stop a service, or change account access immediately.


Before a consequential command:

1. Identify the current user and current directory.
2. Expand variables mentally or print them safely.
3. List the target with a read-only command.
4. Check whether the command overwrites, truncates, removes, or recursively changes data.
5. Create a suitable backup or rollback point when the change is not easily reversible.
6. Run the smallest command that performs the intended action.

7. Verify the result and the exit status.

Chapter 2: The Shell and the Anatomy of a Command

2-14 Safe Execution When There Is No Universal Undo



No universal undo.

017

Figure 2-10: Safe Execution When There Is No Universal Undo

For example, before removing generated files with a wildcard, first print the matches:

```
printf '%s\n' ./*.tmp
```

Only after confirming the displayed targets should you consider a removal command. **Inspection is part of execution, not an optional extra step.**

A Practical Command-Discovery Sequence

Suppose you are told to inspect a service with a command named `systemctl`, but you have not used the host before. With an ordinary account:

```
command -v systemctl
type systemctl
systemctl --help | head
man systemctl
```

The sequence determines whether the command exists, how the shell resolves it, what its concise usage looks like, and where to find full documentation. You can then run a read-only operation and capture its result:

```
systemctl is-system-running
status=$?
printf 'exit status: %s\n' "$status"
```

The text and exit status vary with system state. A host can report `running`, `degraded`, or another documented state. The important practice is to **identify the command, read its interface, run a bounded check, and evaluate both output and status.**