# Introduction to systemd Basics

systemd is the default system and service manager for SUSE Linux Enterprise Server that initializes and manages system components, background daemons and user space processes. By using declarative unit files and target groupings, systemd parallelizes boot operations, manages service dependencies, tracks resources and provides unified logging, enabling administrators to efficiently maintain system stability and streamline service control.

**Prerequisites**

- Basic understanding of Linux processes, daemons and control groups

Publication Date: 05 Nov 2026

## Contents

1 What is systemd? 3

2 About the systemd boot process 3

3 Structure of a unit file 7

4 Unit file types 8

5 Unit dependencies and order 10

6 Logging 10

7 Understanding systemd targets 11

8 Using systemd as a regular user 13

9 **systemctl** commands overview 14

10 systemd troubleshooting 16

11 systemd best practices 17

12 Legal Notice 18

A GNU Free Documentation License 19

# 1 What is systemd?

systemd is a system and service manager for Linux operating systems. It is the default initialization system for major Linux distributions. systemd is not directly initiated by the user, but installed through the /sbin/init and started during the early boot. systemd acts as the init system that brings up and maintains user space services when run as the first process on boot (PID 1). PID 1 is known as *init* and is the first Linux user-mode process created. It runs until the system shutdown.

systemd owns PID 1, and is started directly by the kernel. All other processes are started directly by systemd or one of its child processes. systemd mounts the host's file system and manages temporary files. It is backward compatible with the SysV init scripts. SysV is an initialization system that predates systemd.

In systemd, a unit is a resource that the system knows how to operate on and manage. This is the primary object that the systemd tools use. These resources are defined with configuration files called unit files.

**systemctl** is the central management tool for controlling the init system. It is used to examine and control the state of the systemd system and service manager.

Targets in systemd are groups of related units that act as synchronization points during a system boot. Target unit files have a .target file extension. Target units group together various systemd units through a chain of dependencies.

For troubleshooting, you can use **journalctl**, which is used to query and display log messages from the systemd journal.

For more information on systemd, you can refer to https://systemd.io and **man 1 systemd**.

# 2 About the systemd boot process

The first step in the boot process is to load the Linux kernel, which is the main component of the Linux operating system. Once the kernel is loaded, it initializes the hardware and starts the systemd process, which is the first process that runs on the system.

## 2.1 Linux boot process

The Linux boot process is the initial stage of the operating system's startup. It is the process by which the operating system loads the memory, initializes the components and prepares to execute user applications.

The Linux boot process is divided into four main stages:

**Stage 1: BIOS**

When you power on your computer, your computer starts BIOS (Basic Input/Output System) and performs a POST (Power On Self Test). This is an integrity check that probes the hardware functionality of components such as hard disks, SSD, keyboard, RAM, USB ports and any other hardware. If the hardware works as expected, the boot process moves on to the next stage.

**Stage 2: The boot loader**

Once POST is complete, BIOS searches for and loads the boot loader program stored in the MBR (Master Boot Record). The MBR is a 512-byte code that is usually located at `/dev/sda` or `/dev/hda` depending on your hard drive architecture. The MBR can also be located on a live USB or DVD installation of Linux. BIOS loads and executes this MBR code.

There are three main boot loaders in Linux: LILO, GRUB and GRUB2. The GRUB2 (Grand Unified Bootloader) boot loader is the latest and primary boot loader in modern Linux distributions. The GRUB2 configuration file is located at `/boot/grub2/grub2.cfg`. Once BIOS locates the GRUB2 boot loader, It executes and loads it into the main memory (RAM).

**Stage 3: Linux kernel initialization**

The Linux kernel is the heart of the operating system. In your Linux system, the kernel interfaces with the hardware, controls memory management and manages processes. The boot loader loads the selected Linux kernel. The kernel self extracts from a compressed version and mounts the root file system. It then runs the `/sbin/init` program.

**Stage 4: `systemd`**

The kernel loads `systemd`, which is a system and service manager for Linux operating systems. `systemd` then runs all the other initialization processes.

## 2.2 Boot process with `systemd`

Once the kernel loads `systemd`, `systemd` takes over and starts the other system services that are required to bring the system up and running. This includes services such as networking service, the login manager, etc.

The boot process is parallelized in the order in which specific target units are executed. `systemd` uses the `/etc/systemd/system/default.target` file to determine the target that the Linux system should boot into. This file is a link to `graphical.target` which boots the graphical login manager. `systemd`

activates all target units that are dependencies of `default.target` as well as recursively all dependencies of these dependencies. Once all the services are started, your system is ready to use and the login manager displays. You can now log in and start using the system.

## 2.3 Analyzing the system boot process performance with the **`systemd-analyze`** command

Use the **`systemd-analyze`** command to analyze the performance of the system boot process. The command can also be used to retrieve other state and tracing information from the system and service manager. It is used to check that unit files are correct and also to access special functions useful for advanced system manager debugging.

Some examples include:

***Viewing the time it takes for the system to boot***

```
> systemd-analyze time
Startup finished in 3.404s (kernel) + 2.415s (initrd) + 13.125s (userspace) =
 18.945s
graphical.target reached after 13.117s in userspace
```

***Getting a high-level overview of the boot process, which includes services that are started and the time it takes for each service to start***

```
> systemd-analyze critical-chain
  The time when unit became active or started is printed after the "@" character.
  The time the unit took to start is printed after the "+" character.

  graphical.target @13.117s
  └─multi-user.target @13.117s
    └─getty.target @13.117s
      └─getty@tty1.service @13.116s
        └─plymouth-quit-wait.service @10.775s +2.338s
          └─systemd-user-sessions.service @10.769s +3ms
            └─remote-fs.target @10.764s
              └─iscsi.service @10.747s +16ms
                └─network-online.target @10.744s
                  └─NetworkManager-wait-online.service @1.547s +9.197s
                    └─NetworkManager.service @1.507s +37ms
                      └─network-pre.target @1.504s
                        └─wpa_supplicant.service @2.341s +5ms
                          └─dbus.service @1.042s
                            └─basic.target @1.036s
```

```
                                  └─sockets.target @1.036s
                                    └─snapd.socket @1.035s +590us
                                      └─sysinit.target @1.030s
                                        └─systemd-update-utmp.service @1.025s +5ms
                                          └─auditd.service @976ms +47ms
                                            └─systemd-tmpfiles-setup.service @964ms +9ms
                                              └─local-fs.target @962ms
                                                └─snapd.mounts.target @961ms
                                                  └─snap-core18-2796.mount @417ms +543ms
                                                    └─dev-loop9.device @961ms +628us
```

This command prints a tree of time-critical units either for each of the specified units or for the default target. The initialization of services might depend on socket activation and the parallel execution of units. Similar to the **`blame`** command, it displays the time taken by a unit to activate, which is not defined for units like device units that transition directly to active state.

***Viewing a list of services started during the boot process and displayed according to the time taken by each service***

```
> systemd-analyze blame
  9.197s NetworkManager-wait-online.service
  4.002s fwupd.service
  2.338s plymouth-quit-wait.service
  1.282s dracut-pre-udev.service
  1.062s sys-devices-platform-serial8250-tty-ttyS0.device
  1.062s dev-ttyS0.device
  1.061s dev-ttyS1.device
  1.061s sys-devices-platform-serial8250-tty-ttyS1.device
  1.060s dev-ttyS11.device
  1.060s sys-devices-platform-serial8250-tty-ttyS11.device
  1.059s sys-devices-platform-serial8250-tty-ttyS13.device
  1.059s dev-ttyS13.device
  1.059s sys-devices-platform-serial8250-tty-ttyS10.device
  1.059s dev-ttyS10.device
  1.058s sys-devices-platform-serial8250-tty-ttyS14.device
  1.058s dev-ttyS14.device
  1.058s dev-ttyS12.device
  1.058s sys-devices-platform-serial8250-tty-ttyS12.device
  1.056s sys-devices-platform-serial8250-tty-ttyS17.device
```

The initialization of one service might be slow because it is waiting for another service initialization to complete. It displays the time taken by a unit to activate, which is not defined for units like device units that transition directly to active state. This command does not display results for services with `Type=simple` because `systemd` considers these services to be started immediately, hence the initialization delays cannot be analyzed.

*Generating a vector graphics file that displays the events that take place during the boot process*

```
> systemd-analyze plot > /temp/sample.svg
```

This command creates a SVG file in the `temp` directory. The SVG file is a text file that defines a set of graphics vectors that applications such as LibreOffice Draw use to generate a graph.

# 3 Structure of a unit file

In `systemd`, a unit refers to any resource that the system knows how to operate on and manage. This is the primary object that the `systemd` tools use. These resources are defined using configuration files called unit files. Administration is easier when you understand unit files when working with `systemd`. Unit files use a simple declarative syntax that allows you to see easily the purpose and effects of a unit upon activation. Unit files have sections with directives, for example:

```
[Section]
Directive1=value
Directive2=value
. . .
```

Unit file types include the following sections:

`[Unit]`

The first section found in most unit files is the `[Unit]` section. This section is used to define the metadata of the unit file and configure the relationship of the unit file to other unit files. This section is usually placed at the top because it provides an overview of the unit file.

`[Automount] / [Mount] / [Path] / [Service] / [Slice] / [Socket] /[Swap] / [Timer]`

Sections containing directives that are specific to the respective type. See *Section 4, "Unit file types"* for a list of available types. Note that the types `device`, `target`, `snapshot` and `scope` do not have a type-specific section.

`[Install]`

This is often the last section in the unit file and is optional. This section is used to define the behavior of a unit file when it is enabled or disabled. When you enable a unit file, it automatically starts at boot. Based on the specific unit, there could be a dependency on other related units to work properly. For example, `chrony` requires the directives `After`, `Wants`, and `Before`, which are all dependencies for `chrony` to work with.

EXAMPLE 1: A `systemd` SERVICE FILE

```
[Unit]
Description=usbguard ❶

[Service]
ExecStart=/usr/sbin/usb-daemon ❷

[Install]
WantedBy=multi-user.target ❸
```

❶ A brief and meaningful description explaining the service file's purpose.

❷ Specifies the program to be executed when the service is started.

❸ Starts a multi-user system with networking, and no graphical environment. This directive allows you to specify a dependency relationship.

# 4 Unit file types

You can determine the type of unit by its file extension. `systemd` categorizes units according to the type of resource they describe.

Types of unit files available for `systemd`:

`.service`
: Describes how to manage a service or application. This includes how to start or stop the service, reload its configuration file (if applicable), under what conditions the service starts automatically, and the dependency or the hierarchy information for related unit files.

`.scope`
: This unit file is automatically created by `systemd` from the information received from the D-Bus interface and is used to manage sets of system processes that are created externally.

`.path`
: Defines a path for path-based activation. By default, a `.service` unit file of the same base name is activated. `inotify` is a kernel API that is used by programs that want to be notified about changes to files.

`.snapshot`

The **`systemctl snapshot`** command automatically creates a `.snapshot` unit file. This command creates temporary snapshots of the current state of the system. You can modify the current state of the system after making changes. Snapshots are used for rolling back temporary states.

`.timer`

Defines a timer that is managed by `systemd`. This is similar to a cron job for delayed or scheduled activation. A unit file with the same name, but with file extension `.service` is started when the timer is reached.

`.slice`

Associate Linux Control Group nodes, which allow resources to be assigned or restricted to any processes associated with the slice. The name indicates the hierarchy within the control group tree. Units are placed in slices by default depending on their type.

`.target`

Provides synchronization for other units during a boot up or a change in state, or brings the system to a new state. Other units specify their relation to targets in order to sync with the target's operations.

`.socket`

Describes a network, an IPC socket, or a FIFO buffer that `systemd` uses for socket-based activation. There is an associated `.service` file that starts when an activity is seen on the socket that this unit defines.

`.device`

Defines a device that has been designated for `systemd` management by `udev` or `sysfs` file system. Not all devices have the `.device` file. This unit file is required when ordering, mounting, or accessing a device.

`.swap`

Defines the swap space on the system. The name of the unit file must reflect the device or file path of the space.

`.mount`

Defines a mount point on the system to be managed by `systemd`. This file is named after the mount path, with the slashes changed to dashes. Entries within `/etc/fstab` can have units created automatically.

`.automount`

Defines a mount point that is automatically mounted. Name the file after the mount point that it refers to. A matching `.mount` unit file is required to define the specifics of the mount.

# 5 Unit dependencies and order

`systemd` has two types of dependencies: requirement and order dependencies. Requirement dependencies specify which other units must be either started or stopped when activating a unit. Order dependencies specify the order in which units must be started.

*Unit dependencies*

Unit files have the dependencies feature. A unit may want or require one or more other units before it can run. These dependencies are set in unit files with the directives `Wants` and `Requires`.

`Wants`
: For example, if unit A has `Wants=unit B`, when unit A is run, unit B runs as well. But if unit B starts successfully or not, does not have an influence on unit A running successfully.

`Requires`
: If unit A has `Requires=unit B`, both units run but if unit B does not run successfully, unit A is deactivated. It does not matter if the processes of unit A would have run successfully.

*Unit order*

Without proper instructions, `systemd` can run a group of units at the same time. Starting services in the right order is important for a good functioning of the Linux system. You can arrange the order with the unit file directives `Before` and `After`.

`Before`
: For example, if unit A has `Before=unit B`, when both units are run, unit A is executed fully before unit B.

`After`
: If unit A has `After=unit B`, when both units are run, unit B is executed fully before unit A.

# 6 Logging

Log files and journals are important for system administration. They give in-depth information about a system and are very important for troubleshooting and auditing. Log files contain events and messages generated by the kernel, applications, and users that log into the system. You can use the **`journalctl`** command to query the journal. This command views logs collected by `systemd`. The `systemd-journald` service handles `systemd`'s log collection. `systemd-journald` saves the events and messages in a binary format.

# 7 Understanding systemd targets

systemd uses units and targets. A systemd unit defines a service or action on the system, which consists of a name, type, and configuration file.

## 7.1 What are systemd targets?

A systemd target combines several units and defines which services have to be started to reach the target. On a server, for example, this is a state where the network is running and multiple users can log in. These files are identified by the suffix .target.

Similar to unit files, different targets may be nested via dependencies. For example, multi-user.target requires (among others) the targets that set up login and user session services.

Common systemd targets:

default.target
: Boots by default. The default.target file is a symbolic link to the true target file, such as graphical.target for a desktop workstation. For a server, it is usually graphical.target.

poweroff.target
: Shuts down and powers off the system.

rescue.target
: Target unit that pulls the base system and starts a rescue shell session.

multi-user.target
: Sets up a non-graphical (console) multi-user system.

graphical.target
: Uses a graphical multi-user system with network services.

reboot.target
: Shuts down and reboots the system.

## 7.2 Setting up a systemd target

To set up a systemd target, you need to create a new target unit file and define any required dependencies. A target is a group of unit files that manages the state of services. A target can be used to control groups of services or to define system states.

1. Create a new file with the `.target` extension in `/etc/systemd/system/`.
   For example:

```
> sudo  vi /etc/systemd/system/test.target
```

2. Define the target with basic configuration.
   For example:

```
[Unit]
Description= Test target
Requires=multi-user.target
After=multi-user.target

[Install]
WantedBy=multi-user.target
```

   - *Requires:* specifies that `multi-user.target` must be reached for `test.target` to be active.
   - *After:* specifies that this target starts after the `multi-user.target`.
   - *WantedBy:* specifies that this target should be part of `multi-user.target`.

3. Set permissions for the target file.
   For example:

```
> sudo  chmod 644 /etc/systemd/system/test.target
```

4. Link the required services to the target file. Create a `.wants` directory and then create symbolic links to the services you want to include in your custom target.
   For example:

```
> sudo  mkdir -p /etc/systemd/system/custom.target.wants/
```

```
> sudo  ln -s /etc/systemd/system/test.service /etc/systemd/system/
custom.target.wants/
```

5. Reload `systemd`.
   For example:

```
> sudo  systemctl daemon-reload
```

6. Start and enable the target.
   For example:

```
> sudo  systemctl start test.target
```

```
> sudo  systemctl enable test.target
```

7. Verify the status of the custom target.
   For example:

```
> sudo  systemctl status test.target
```

8. Verify all the units that are part of this target.
   For example:

```
> sudo  systemctl list-dependencies test.target
```

9. Your custom `systemd` target is set up and configured.

You can use the following common configuration options:

- *Description:* A human-readable description of the target.
- *Documentation:* URIs pointing to documentation.
- *AllowIsolate:* If set to `true`, this unit can be used with the **`systemctl isolate`** command.
- *Wants:* Units that can be started but are not mandatory.
- *After/Before:* Defines the order of dependencies.
- *Conflicts:* Units that should not be active while this target is active.

For more information on `systemd` targets, refer to *man 5 systemd.target* and *man 7 systemd.special*.

# 8 Using `systemd` as a regular user

You can use `systemd` as a regular user for better security or when you do not have `root` user privileges. Running an unprivileged service can be done by creating a `user` service.

When creating and using a user service, consider the following:

- User service sessions are terminated when the user's session ends. This can be overridden by using the **`loginctl enable-linger`** *`USERNAME`* command.
- User service files are located in `/etc/systemd/user` or `$HOME/.config/systemd/user/`.
- You can control user services with the **`systemctl --user`** command.

# 9 `systemctl` commands overview

The **`systemctl`** command is used to examine and control the state of `systemd` and service manager.

You can use the following common **`systemctl`** commands and refer to the *man systemctl* page.

## 9.1 Viewing `systemd` information

To view information about `systemd` components, you can use the following commands:

**systemctl list-units**
: Lists the `systemd` units. You can use the optional arguments: **`--state=running`** to show the active units and **`--type=service`** to show the exited and active units.

**systemctl list-unit-files**
: Lists the `systemd` units and the status, such as static, generated, disabled, alias, masked, and enabled.

**systemctl list-dependencies**
: Lists the dependency tree.

**systemctl list-dependencies** *`UNIT_FILE`*
: Lists the dependencies of a unit file.

## 9.2 Managing `systemd` services

The **`systemctl`** command enables you to perform the following tasks with services.

**systemctl status** *`SERVICE`*
: Checks the status of the specific service.

**systemctl show** *SERVICE*

Displays the service information.

**systemctl start** *SERVICE*

Instead of manually starting the service, use the **`start`** command. When a change is made to the configuration file, the related service must be started again.

**systemctl stop** *SERVICE*

Stops a specific running service.

**systemctl restart** *SERVICE*

Instead of manually restarting the service, use the **`restart`** command. When a change is made to the configuration file, the related service must be restarted again.

**systemctl enable** *SERVICE*

Enables the service on boot.

**systemctl disable** *SERVICE*

Disables the service on boot.

**systemctl reload-or-restart** *SERVICE*

Reload the service if it supports reloading, otherwise it restarts the service. If the service is not running, it is restarted.

**systemctl mask** *SERVICE*

When a service is masked, this means the unit file is symlinked to `/dev/null`. A symlink for a masked service is created from `/etc/systemd/system` to point to `/dev/null`. This makes it impossible to load the service even if another enabled service requires it. It must be stopped manually, or it continues to run in the background. You can use **`--runtime`** option to only mask temporarily until the next reboot of the system.

```
Created symlink /etc/systemd/system/FOSSLinux.service → /dev/null.
```

**systemctl unmask** *SERVICE*

Unmasks the service. It is effective when the system is started or restarted manually.

## 9.3 Managing system states

The **`systemctl`** command enables you to perform power management processes on your system, like restarting, shutting down and so on, as described below.

**systemctl reboot**

Reboots the system `reboot.target`.

**systemctl poweroff**

Powers off the system `poweroff.target`.

**systemctl emergency**

Goes into the emergency mode `emergency.target`.

**systemctl default**

Goes back to default target `multi-user.target`.

# 10 systemd troubleshooting

You can use the following troubleshooting tips to identify and resolve issues with `systemd` services and ensure a smooth system operation.

***Check the syntax of your `systemd` unit file with the `systemd-analyze verify SERVICE`***

Before starting or enabling a `systemd` service, check the syntax of the unit file to ensure there are no errors. For example:

```
> sudo systemd-analyze verify /etc/systemd/system/my-custom-service.service
```

The command analyzes the unit file and reports any syntax errors, missing files, or other issues. You must fix any reported issues before enabling and starting the service.

***Check the logs for your service with the `journalctl -u SERVICE` command***

If you experience any issue with a `systemd` service, check the service's log. For example:

```
> sudo  journalctl -u my-custom-service.service
```

The command displays logs for the specified service, including any error messages, warnings, or other relevant information. You can use these logs to identify and fix issues with the service.

***Use the `systemd-analyze plot` command to visualize the boot process***

If a service is causing issues during the boot process, you can use the **`systemd-analyze plot` command** to visualize the boot process and identify issues. For example:

```
> sudo systemd-analyze plot > boot-plot.svg
```

The command creates an SVG file called `boot-plot.svg` that contains a graphical representation of the boot process and potential issues. This includes the start and stop time of each service. You can open this file in an SVG-compatible image viewer or Web browser to analyze services that are causing issues during the boot process.

#### *Troubleshoot failed services*

To find out which services have failed and to inspect the log output:

```
> sudo systemctl --state=failed
```

#### *Check the runtime status of a service*

To find out the current runtime status of a service:

```
> sudo systemctl status SERVICE
```

#### *Shutdown or reboot takes long*

If the shutdown or reboot takes long, it could be a service that is not exiting. `systemd` waits for some time for each service to exit before trying to terminate it. A common issue is a suspended service or a stalled shutdown. To find out, use the following:

```
> sudo systemctl poweroff
  Failed to power off system via logind: There's already a shutdown or sleep
 operation in progress
```

```
> sudo systemctl list-jobs
```

You can cancel the running and waiting jobs, and again shut down or reboot:

```
> sudo systemctl cancel
```

```
> sudo systemctl stop systemd-suspend.service
```

## 11 `systemd` best practices

You can follow some of the best practices to ensure efficient `systemd` services that are equipped to handle different situations.

#### *Check the runtime status of a service*

To find out the current runtime status of a service:

```
> sudo systemctl status SERVICE
```

#### *Use absolute path in your `systemd` unit file*

Use an absolute path for executable files and required files, such as configuration files or scripts in your `systemd` unit file. `systemd` does not rely on the user's environment variables like `$PATH` to locate files.

#### *Use the ExecReload directive*

Use the *ExecReload* directive in the `[SERVICE]` section when you want to define a specific command that should be executed when you reload a service with the **`systemctl reload`** command. This is useful for services that can dynamically reload their configuration without a restart.

```
[Service]
ExecStart=PATH_TO_EXECUTABLE
ExecReload=PATH_TO_RELOAD_SCRIPT
```

#### *Use the RestartSec directive*

Use the *RestartSec* directive in the `[SERVICE]` section when you want to define a delay (in seconds) before the service is restarted after a failure. This is useful for services that require a specified time to release resources or prevent rapid restart loops that can cause high system load.

```
[Service]
ExecStart=PATH_TO_EXECUTABLE
Restart=on-failure
RestartSec=5
```

#### *Disable emergency mode on a remote machine*

You can disable emergency mode on a remote machine, for example, a virtual machine hosted on Google Cloud. If this mode is enabled, the machine is blocked from connecting to the network. For example:

```
> sudo systemctl mask emergency.service
```

```
> sudo systemctl mask emergency.target
```

# 12 Legal Notice

Copyright© 2006– 2026 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled “GNU Free Documentation License”.

For SUSE trademarks, see https://www.suse.com/company/legal/. All other third-party trademarks are the property of their respective owners. Trademark symbols (®, ™ etc.) denote trademarks of SUSE and its affiliates. Asterisks (*) denote third-party trademarks.

All information found in this book has been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. Neither SUSE LLC, its affiliates, the authors, nor the translators shall be held liable for possible errors or the consequences thereof.

# A GNU Free Documentation License

Copyright (C) 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

### 0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondarily, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

### 1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The

"Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

## 2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or non-commercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

## 3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

## 4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.

C. State on the Title page the name of the publisher of the Modified Version, as the publisher.

D. Preserve all the copyright notices of the Document.

E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.

F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.

G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.

H. Include an unaltered copy of this License.

I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.

J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.

K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.

L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.

M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.

N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.

O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties--for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

## 5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

## 6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

## 7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

## 8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

## 9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

## 10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See https://www.gnu.org/copyleft/.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

## ADDENDUM: How to use this License for your documents

```
Copyright (c) YEAR YOUR NAME.
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.2
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.
A copy of the license is included in the section entitled “GNU
Free Documentation License”.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.