Kdump Configuration and Troubleshooting Guide

Kdump is a Linux kernel crash dumping mechanism in SUSE Linux Enterprise Server for NVIDIA that captures a snapshot of physical memory (`vmcore`) when the kernel encounters a fatal error. By leveraging Kexec to instantly boot a secondary capture kernel in reserved RAM, Kdump bypasses corrupted primary kernel states. This enables administrators and developers to reliably perform root-cause analysis and debug silent system crashes without relying on failing logging daemons.

Prerequisites

- Linux command-line fundamentals
- Adequate RAM reserved for the crash kernel via the crashkernel boot parameter

Publication Date: 05 Nov 2026

Contents

1 What is Kdump and how does it work? 3

2 Installing and configuring Kdump 5

3 How to troubleshoot common Kdump issues 10

4 For more information 13

5 Legal Notice 13

A GNU Free Documentation License 13

1 What is Kdump and how does it work?

Kdump is a kernel crash dumping mechanism that captures the system’s memory state into a vmcore file when a system crash occurs. A vmcore file is a snapshot of your computer's system memory (RAM) taken at the exact moment the Linux kernel crashed.

1.1 Why is Kdump critical for system crash debugging?

The primary importance of Kdump lies in its ability to capture a snapshot of a system's memory at the exact moment of a critical failure.

When a Linux kernel experiences a fatal error, syslog or journald usually fail along with it, often leaving no record of what went wrong. Kdump bypasses this limitation by using Kexec to boot a secondary capture kernel in a reserved slice of RAM. The unstable crashed system is replaced with a freshly started and stable capture kernel.

Without this tool, administrators are often left with nothing but a blank screen or a frozen console, making it nearly impossible to diagnose the root cause of intermittent or silent system crashes.

1.2 Understanding the dual-kernel model

Kdump uses a second isolated kernel referred to as the capture or crash kernel to handle a system crash safely. When the main system kernel fails, you cannot trust it to write its own crash logs to disk— because the kernel memory might be corrupted and the kernel itself is no longer reliable. The dual-kernel approach solves this by jumping into a completely different environment.

The model relies on two distinct kernels residing in memory simultaneously:

- **The production (primary) kernel** is the kernel you use every day. It runs your applications and services.
- **The capture (crash) kernel** is a second copy of the kernel loaded in a small reserved area of RAM. It only starts when the primary kernel panics. Alongside the crash kernel, a special-purpose `initramfs` image is loaded to the reserved RAM. It is built by the Kdump tool and includes all the drivers, settings, and programs to store the vmcore file.

1.3 What is Kexec and how does it boot the crash kernel?

Kexec is a system call that functions as a software-defined boot loader, allowing a running kernel to bypass the hardware BIOS/UEFI stage and directly hand over control to a new kernel. By loading the secondary kernel's image and parameters into memory while the system is still active, Kexec performs a warm boot that preserves the state of RAM and significantly reduces downtime. This mechanism is the backbone of the Kdump dual-kernel model, as it provides a reliable way to jump from a crashing production environment into a clean recovery environment for data capture.

The most important component of Kexec is the **`kexec`** command. You can load a kernel with Kexec in two ways:

- Load the kernel to the address space of a production kernel for a regular reboot:

  ```
  > sudo  kexec -l KERNEL_IMAGE
  ```

  You can later boot to this kernel with the command **`kexec -e`**. Instead of using Kexec, you can directly use a program called `kexec-bootloader`. It finds the default kernel, initrd and command-line options from the boot loader configuration and passes everything to Kexec to load the default kernel properly.

- Load the kernel to a reserved area of memory:

  ```
  > sudo  kexec -p KERNEL_IMAGE
  ```

  This kernel is booted automatically when the system crashes.This is what Kdump uses to load the capture kernel.

1.4 What data is captured in a Linux vmcore file?

A vmcore file is a snapshot of your system's physical memory (RAM) taken at the exact moment the Linux kernel crashed. When kdump is set up on a system, the kdump service loads the capture kernel and `initramfs` on system boot, using the `kexec` program, into a pre-reserved area of RAM— the crash kernel area. If at some point the system crashes, the capture kernel is started. Its memory is restricted to the pre-reserved crash kernel area, so none of the memory used by the crashed production kernel is overwritten. The job of the capture kernel and `initramfs` is to save the contents of the production kernel's memory into a vmcore file.

The vmcore file is a snapshot of the RAM and includes:

- **The Kernel state:** All active kernel data structures, global variables and the call stack, showing what the CPU was executing at the time of the crash.
- **Process information:** A list of every process that was running, including their individual stacks and registers.
- **Memory pages:** Depending on your settings, it can contain the actual data held in RAM by applications.
- **VMCOREINFO:** special section that tells analysis tools how the kernel's memory was laid out so they can make sense of the raw data.

2 Installing and configuring Kdump

This section describes how to install and configure Kdump.

To install Kdump, run the following command:

This command downloads the following packages:

- **`kdump`**
- **`kexec-tools`**
- **`makedumpfile`**

To store the vmcore file remotely, you can install any of the following recommended packages:

- Install the `cifs-utils` package to store on Samba stores.
- Install the `lftp` package to store on FTP / SFTP servers.
- Install the `nfs-client` package to store on NFS servers.
- Install the `openssh-clients` package to store over SSH.



Important: Managing Kdump and GRUB configuration conflicts

The `kdump-commandline` service automatically manages your crash memory allocation. It will overwrite any manual `crashkernel=` entries in `GRUB_CMDLINE_LINUX_DEFAULT` inside `/etc/default/grub` unless the automation is explicitly disabled. To prevent this behavior, modify the `KDUMP_UPDATE_BOOTLOADER` variable to `false`.

How do you stop Kdump from overwriting GRUB configurations?

To keep your manual `crashkernel=` settings from being erased, update your configuration with these rules:

- Set `KDUMP_UPDATE_BOOTLOADER="false"` to disable automatic updates.
- Keep all `crashkernel=` parameters out of the `GRUB_CMDLINE_LINUX` line.
- Manual configurations belong exclusively in `GRUB_CMDLINE_LINUX_DEFAULT`.

Why should you avoid adding crash kernel settings to `GRUB_CMDLINE_LINUX`?

Adding `crashkernel=` to `GRUB_CMDLINE_LINUX` causes system conflicts.

- Settings in `GRUB_CMDLINE_LINUX` apply globally to all boot targets.
- This includes the rescue system, which does not run Kdump.
- The `kdump-commandline` service never alters this specific line.
- Leaving values here creates direct conflicts with automated service settings.

2.1 Configuring Kdump for non-immutable systems

To use Kdump, a specific amount of system memory (RAM)—often referred to as crash kernel memory-must be permanently reserved at boot time for the capture kernel and initramfs. The amount of memory to reserve is specified using the `crashkernel= kernel` command-line parameter. Kdump includes `kdumptool calibrate` and `kdumptool commandline` to calculate exactly how much memory needs to be reserved.

When the **`systemctl enable --now kdump`** command triggers the main `kdump.service`, it automatically calls a helper service named `kdump-commandline.service` before attempting to load the kernel. This helper service verifies if the current kernel command line contains the expected `crashkernel=` options. The options are:

- Fixed memory allocation (e.g. ,`256M`) reserves a strict hardcoded amount of RAM.
- Range-based allocation (e.g. `2G-8G:192M`) dynamically scales the reserved memory based on the total RAM installed in the system.

- Automatic allocation (`auto`) where the kernel automatically calculates the required memory based on system architecture.
- High/low split (e.g. `512M,high crashkernel=128M,low`) allocates memory across different addressable zones specifically to accommodate hardware drivers on large enterprise systems.

If the options are missing, the helper service automatically modifies the boot loader configuration to contain the crash kernel options to be applied on the next boot.

Before attempting to load, the `kdump.service`:

- Starts the `kdump-commandline.service`
- Updates the boot loader
- Builds the capture initramfs

To verify that the service successfully checked and updated your kernel boot options, inspect the system logs:

```
> sudo  journalctl -u kdump-commandline.service
```

Expected output:

```
  Started Check and update kdump options on the kernel command line.
  Kdump expects these crashkernel= values on the kernel command line:
  crashkernel=72M,low crashkernel=443M,high
  (based on the output of 'kdumptool calibrate')
  No crashkernel= found in /proc/cmdline. Kdump can't be started.
  Updated the command line options in the bootloader.
  Changes will take effect on next boot.
```

The boot loader has been updated. Reboot the machine to apply the changes and start Kdump.

Other configurations:

The `/etc/sysconfig/kdump` file is the primary configuration file for managing the Kdump service. It includes essential options such as:

- `KDUMP_CRASHKERNEL` for defining the reserved crash memory size
- `KDUMP_VERBOSE` for increasing debugging log levels
- `KDUMP_SAVEDIR` for specifying the core dump (vmcore) storage location

EXAMPLE 1: CONTENTS OF A `/etc/sysconfig/kdump` FILE

```
cat /etc/sysconfig/kdump
  ## Path:        System/Kernel/Kdump
```

```
## Description: Crash Dump Configuration

## Type:        string
## Default:     "auto"
## ServiceRestart:      kdump
#
# KDUMP_CRASHKERNEL is used by "kdumptool commandline" to determine
# the value of the kernel "crashkernel=" command-line options.
#
# Possible values are "auto" or one or more "crashkernel=..." values.
#
# KDUMP_CRASHKERNEL="auto" makes "kdumptool commandline" use the default
# values as proposed by "kdumptool calibrate".
#
# If the default values are not adequate you may provide a manual setting,
# e.g., KDUMP_CRASHKERNEL="crashkernel=72M,low crashkernel=300M,high"
#
# See also: kdump(5).
#
KDUMP_CRASHKERNEL="auto"


## Type:        boolean
## Default:     "true"
## ServiceRestart:      kdump
#
# When KDUMP_UPDATE_BOOTLOADER is set to "true" "kdumptool commandline -c -U"
# called from the kdump-commandline.service  will check if the expected kernel
# command-line options are present and update the bootloader using *pbl*(8) if
# not.
#
# Also the kdump options are removed from the command line when kdump is
# uninstalled.
#
# When set to false "kdumptool commandline -c -U" will only check the options
# and report possible warnings but bootloader will not be updated.
#
KDUMP_UPDATE_BOOTLOADER="true"

## Type:        string
## Default:     ""
## ServiceRestart:      kdump
#
# Kernel Version string for the -kdump kernel. Use an empty string to
# auto-detect a suitable kernel.
#
# See also: kdump(5).
```

```
  #
  KDUMP_KERNELVER=""

  ## Type:        integer
  ## Default:     32
  ## ServiceRestart:      kdump
  #
  # Number of CPUs to be used in the kdump environment. You may want to
  # increase the number if computing power is the bottleneck in your setup.
  #
  # If the value is zero, use all available CPUs.
  #
  # Decreasing the number of CPUs lowers the memory required by kdump.
  # Increasing it raises memory usage and may cause failures if crashkernel
  # memory is insufficient.
  #
  # See also: kdump(5).
  #
  KDUMP_CPUS=32

  ## Type:        string
  ## Default:     ""
  ## ServiceRestart:      kdump
  #
  # The kdump commandline is the command line that needs to be passed off to
  # the kdump kernel. If a command line is not specified, the default will be
  # taken from /proc/cmdline and adjusted.
  #
  # WARNING: There are a few options that always should be passed to the kdump
  # kernel. See kdump(5) for details. Don't use that variable, consider using
  # KDUMP_COMMANDLINE_APPEND instead.
  #
  KDUMP_COMMANDLINE=""

  ## Type:        string
  ## Default:     ""
  ## ServiceRestart:     kdump
  #
  # Set this variable if you only want to _append_ values to the default
  # command line string. The string also gets appended if KDUMP_COMMANDLINE
  # is set.
  #
  # For network based dumps, you may have to add a "net_delay" parameter
  # here. Consult the man page for details.
[...]
```

For more information, refer to `man kdump(5)` and `man kdump(7)`.

Any modifications made to this configuration file require a manual restart of the Kdump service to take effect. However, if any changes are made specifically to the `KDUMP_CRASHKERNEL` memory allocation size, a full system reboot is required because these settings alter the system's core boot parameters.

2.2 Testing Kdump

It is advisable to test Kdump after configuring it by simulating a kernel crash. Otherwise, you may only find out that it does not work when an actual kernel crash occurs, leaving you with no possibility to debug the crash. Ensure no critical workloads are running and no unsaved data is present on the system. Additionally, ensure to sync and unmount file systems:

Force an immediate emergency sync of all cached memory data to your hard drives to prevent data loss:

```
echo s > /proc/sysrq-trigger
```

Force an immediate emergency remount of all file systems to read-only mode to prevent data corruption:

```
echo u > /proc/sysrq-trigger
```

Then you can simulate a kernel crash:

```
echo c > /proc/sysrq-trigger
```

Check if there is a new directory created under your `KDUMP_SAVEDIR`, which is `/var/crash` by default. This contains the dmesg and vmcore of the crashed kernel.

3 How to troubleshoot common Kdump issues

Troubleshooting Kdump is a critical process to ensure that your system can successfully capture a vmcore file during a kernel crash, which is often the only way to diagnose system crashes. When troubleshooting Kdump, the process usually fails at one of three stages: during boot (memory reservation), during the crash when the dump does not start, or during the save process when the file is not written.

3.1 Crash kernel memory reserved is insufficient

One of the most common reasons Kdump fails is that the amount of crash kernel memory reserved is insufficient. Different system configurations may require more memory than estimated by **`kdumptool calibrate`** and set up automatically in the boot loader config by the `kdump-commandline.service`.

During Kdump, if you see error messages mentioning low memory and invoking the Out of Memory (OOM) killer, this is the likely cause. If you don't see such messages, trying with increased crash kernel reservation is a good first step.

1. Find the size of the current automatically configured crash kernel reservation by checking the kernel command line:

```
> sudo  cat /proc/cmdline
```

   - `crashkernel=X M`
   - `crashkernel=Y M,low`
   - `crashkernel=Z M,high`

   Add up the values of *X*, *Y*, and *Z* to determine the total size of the current reservation in MiB (*current*).

2. Manually set the new reservation by editing the `/etc/sysconfig/kdump` file. Locate and change the value of the `KDUMP_CRASHKERNEL` variable to double the current size:

```
KDUMP_CRASHKERNEL="crashkernel=<2 * current>M"
```

3. Reboot the system to apply the new memory parameters:

```
> sudo  reboot
```

4. Repeat this process a few times until Kdump successfully executes without errors.



Note

Fine-tuning crash kernel memory involves finding the smallest RAM reservation that successfully captures a crash dump without triggering OOM errors in the capture kernel. Fine-tune your final setting using the range between your last working and non-working values to avoid wasting system memory.

3.1.1 Other troubleshooting issues

Issues with switching to a text virtual terminal

Check the Kdump boot messages for errors. If the screen is blank, the system is likely failing to switch to the text virtual terminal. To fix this, configure a serial console to capture the output; you cannot troubleshoot the issue without these messages.

Once you get some output, you can increase Kdump verbosity. Setting `KDUMP_VERBOSE` to 11 turns on debugging output during all stages of the Kdump process and results in the following behavior:

- Removes the quiet option from the capture kernel command line.
- Runs the `kdump-save` script with -x.
- Runs `makedumpfile` with debugging.

`makedumpfile` or network config errors

During the capture kernel boot, if you see problems with `makedumpfile`, network config, etc., set `KDUMP_CONTINUE_ON_ERROR=false` or `KDUMP_IMMEDIATE_REBOOT=false`. This opens up an interactive shell when an error occurs or when `makedumpfile` finishes, enabling you to manually debug problems from within the capture system.

Dracut errors

If you see dracut errors during the boot of the capture system, add dracut debugging options to `KDUMP_COMMANDLINE_APPEND`.

`kdump initramfs` does not generate correctly

If `kdump initramfs` does not generate correctly, turn on debugging when generating it; **`mkdumprd -fd`**. You can also unpack and inspect the generated initramfs image with the **`lsinitrd --unpack /var/lib/kdump/initrd`** command.

4 For more information

For information on Kdump, refer to the following resources:

- Man pages:
  - man 7 kdump
  - man 5 kdump
- Official Linux kernel documentation: https://docs.kernel.org/admin-guide/kdump/kdump.html
- Man page for Kexec: https://man7.org/linux/man-pages/man8/kexec.8.html

5 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.