

SUSE Linux Enterprise High Availability Extension 15

Installation and Setup Quick Start

This document guides you through the setup of a very basic two-node cluster, using the bootstrap scripts provided by the `ha-cluster-bootstrap` package. This includes the configuration of a virtual IP address as a cluster resource and the use of SBD on shared storage as a node fencing mechanism.

Publication Date: January 05, 2023

Contents

- 1 Usage Scenario 3
- 2 System Requirements 3
- 3 Overview of the Bootstrap Scripts 5
- 4 Installing the High Availability Packages 7
- 5 Using SBD for Node Fencing 7
- 6 Setting Up the First Node 8
- 7 Adding the Second Node 11
- 8 Testing the Cluster 12
- 9 Next Steps 15
- 10 For More Information 15

1 Usage Scenario

The procedures in this document will lead to a minimal setup of a two-node cluster with the following properties:

- Two nodes: `alice` (IP: `192.168.1.1`) and `bob` (IP: `192.168.1.2`), connected to each other via network.
- A floating, virtual IP address (`192.168.2.1`) that allows clients to connect to the service no matter which node it is running on. This IP address is used to connect to the graphical management tool Hawk2.
- A shared storage device, used as SBD fencing mechanism. This avoids split brain scenarios.
- Failover of resources from one node to the other if the active host breaks down (*active/passive* setup).

You can use the two-node cluster for testing purposes or as a minimal cluster configuration that you can extend later on. Before using the cluster in a production environment, modify it according to your requirements.

2 System Requirements

This section informs you about the key system requirements for the scenario described in [Section 1](#). To adjust the cluster for use in a production environment, refer to the full list in *Book "Administration Guide", Chapter 2 "System Requirements and Recommendations"*.

2.1 Hardware Requirements

Servers

Two servers with software as specified in [Section 2.2, "Software Requirements"](#).

The servers can be bare metal or virtual machines. They do not require identical hardware (memory, disk space, etc.), but they must have the same architecture. Cross-platform clusters are not supported.

Communication Channels


At least two TCP/IP communication media per cluster node. The network equipment must support the communication means you want to use for cluster communication: multicast or unicast. The communication media should support a data rate of 100 Mbit/s or higher. For a supported cluster setup two or more redundant communication paths are required. This can be done via:

- Network Device Bonding (preferred).
- A second communication channel in Corosync.

Node fencing/STONITH

A node fencing (STONITH) device to avoid split-brain scenarios. This can be either a physical device (a power switch) or a mechanism like SBD (STONITH by disk) in combination with a watchdog.

This document describes using SBD for node fencing. To use SBD, the following requirements must be met:

- A shared storage device. For information on setting up shared storage, see the *Storage Administration Guide for SUSE Linux Enterprise Server* (<https://documentation.suse.com/sles/html/SLES-all/book-storage.html>) .
- The path to the shared storage device must be persistent and consistent across all nodes in the cluster. Use stable device names such as `/dev/disk/by-id/dm-uuid-part1-mpath-abcdef12345`.
- The SBD device *must not* use host-based RAID, LVM2, nor reside on a DRBD* instance.

For more information on STONITH, see *Book "Administration Guide", Chapter 11 "Fencing and STONITH"*. For more information on SBD, see *Book "Administration Guide", Chapter 12 "Storage Protection and SBD"*.


2.2 Software Requirements

All nodes that will be part of the cluster need at least the following modules and extensions:

- Base System Module 15
- Server Applications Module 15
- SUSE Linux Enterprise High Availability Extension 15

2.3 Other Requirements and Recommendations

Time Synchronization

Cluster nodes must synchronize to an NTP server outside the cluster. Since SUSE Linux Enterprise High Availability Extension 15, chrony is the default implementation of NTP. For more information, see the *Administration Guide for SUSE Linux Enterprise Server 15* (<https://documentation.suse.com/sles/15-GA/html/SLES-all/cha-ntp.html>) .

If nodes are not synchronized, the cluster may not work properly. In addition, log files and cluster reports are very hard to analyze without synchronization. If you use the bootstrap scripts, you will be warned if NTP is not configured yet.

Host Name and IP Address

- Use static IP addresses.
- Only the primary IP address is supported.
- List all cluster nodes in the `/etc/hosts` file with their fully qualified host name and short host name. It is essential that members of the cluster can find each other by name. If the names are not available, internal cluster communication will fail.

SSH

All cluster nodes must be able to access each other via SSH. Tools like `crm report` (for troubleshooting) and Hawk2's *History Explorer* require passwordless SSH access between the nodes, otherwise they can only collect data from the current node.

If you use the bootstrap scripts for setting up the cluster, the SSH keys will automatically be created and copied.

3 Overview of the Bootstrap Scripts

All commands from the `ha-cluster-bootstrap` package execute bootstrap scripts that require only a minimum of time and manual intervention.

- With `ha-cluster-init`, define the basic parameters needed for cluster communication. This leaves you with a running one-node cluster.
- With `ha-cluster-join`, add more nodes to your cluster.
- With `ha-cluster-remove`, remove nodes from your cluster.

All bootstrap scripts log to `/var/log/ha-cluster-bootstrap.log`. Check this file for any details of the bootstrap process. Any options set during the bootstrap process can be modified later with the YaST cluster module. See *Book "Administration Guide", Chapter 4 "Using the YaST Cluster Module"* for details.

Each script comes with a man page covering the range of functions, the script's options, and an overview of the files the script can create and modify.

The bootstrap script **`ha-cluster-init`** checks and configures the following components:

NTP

Checks if NTP is configured to start at boot time. If not, a message appears.

SSH

Creates SSH keys for passwordless login between cluster nodes.

Csync2

Configures Csync2 to replicate configuration files across all nodes in a cluster.

Corosync

Configures the cluster communication system.

SBD/Watchdog

Checks if a watchdog exists and asks you whether to configure SBD as node fencing mechanism.

Virtual Floating IP

Asks you whether to configure a virtual IP address for cluster administration with Hawk2.

Firewall

Opens the ports in the firewall that are needed for cluster communication.

Cluster Name

Defines a name for the cluster, by default `hacluster`. This is optional and mostly useful for Geo clusters. Usually, the cluster name reflects the location and makes it easier to distinguish a site inside a Geo cluster.


QDevice/QNetd

Asks you whether to configure QDevice/QNetd to participate in quorum decisions. We recommend using QDevice and QNetd for clusters with an even number of nodes, and especially for two-node clusters.

This configuration is not covered here, but you can set it up later as described in *Book "Administration Guide", Chapter 13 "QDevice and QNetd"*.

4 Installing the High Availability Packages

The packages for configuring and managing a cluster are included in the High Availability installation pattern. This pattern is only available after the SUSE Linux Enterprise High Availability Extension is installed.

You can register to the SUSE Customer Center and install the High Availability Extension while installing SUSE Linux Enterprise Server, or after installation. For more information, see the *Deployment Guide* (<https://documentation.suse.com/sles/html/SLES-all/cha-register-sle.html>)  for SUSE Linux Enterprise Server.

PROCEDURE 1: INSTALLING THE HIGH AVAILABILITY PATTERN

1. Install the High Availability pattern from the command line:

```
# zypper install -t pattern ha_sles
```

2. Install the High Availability pattern on *all* machines that will be part of your cluster.



Note: Installing software packages on all nodes

For an automated installation of SUSE Linux Enterprise Server 15 and the High Availability Extension, use AutoYaST to clone existing nodes. For more information, see *Book “Administration Guide”, Chapter 3 “Installing the High Availability Extension”, Section 3.2 “Mass Installation and Deployment with AutoYaST”*.

5 Using SBD for Node Fencing

Before you can configure SBD with the bootstrap script, you must enable a watchdog on each node. SUSE Linux Enterprise Server ships with several kernel modules that provide hardware-specific watchdog drivers. The High Availability Extension uses the SBD daemon as the software component that “feeds” the watchdog.

The following procedure uses the softdog watchdog.



Important: Softdog Limitations

The softdog driver assumes that at least one CPU is still running. If all CPUs are stuck, the code in the softdog driver that should reboot the system will never be executed. In contrast, hardware watchdogs keep working even if all CPUs are stuck.

Before using the cluster in a production environment, we highly recommend replacing the `softdog` module with the hardware module that best fits your hardware.

However, if no watchdog matches your hardware, `softdog` can be used as kernel watchdog module.

PROCEDURE 2: ENABLING THE SOFTDOG WATCHDOG FOR SBD

1. On each node, enable the softdog watchdog:

```
# echo softdog > /etc/modules-load.d/watchdog.conf
# systemctl restart systemd-modules-load
```

2. Test if the softdog module is loaded correctly:

```
# lsmod | grep dog
softdog          16384  1
```

6 Setting Up the First Node

Set up the first node with the `ha-cluster-init` script. This requires only a minimum of time and manual intervention.

PROCEDURE 3: SETTING UP THE FIRST NODE (alice) WITH `ha-cluster-init`

1. Log in as `root` to the physical or virtual machine to use as cluster node.
2. Start the bootstrap script:

```
# ha-cluster-init --name CLUSTERNAME
```

Replace the `CLUSTERNAME` placeholder with a meaningful name, like the geographical location of your cluster (for example, `amsterdam`). This is especially helpful to create a Geo cluster later on, as it simplifies the identification of a site.

If you need unicast instead of multicast (the default) for your cluster communication, use the option `-u`. After installation, find the value `udpu` in the file `/etc/corosync/corosync.conf`. If `ha-cluster-init` detects a node running on Amazon Web Services (AWS), the script will use unicast automatically as default for cluster communication.

The script checks for NTP configuration and a hardware watchdog service. It generates the public and private SSH keys used for SSH access and Csync2 synchronization and starts the respective services.

3. Configure the cluster communication layer (Corosync):

- a. Enter a network address to bind to. By default, the script will propose the network address of `eth0`. Alternatively, enter a different network address, for example the address of `bond0`.
- b. Enter a multicast address. The script proposes a random address that you can use as default. Of course, your particular network needs to support this multicast address.
- c. Enter a multicast port. The script proposes `5405` as default.

4. Set up SBD as the node fencing mechanism:

- a. Confirm with `y` that you want to use SBD.
- b. Enter a persistent path to the partition of your block device that you want to use for SBD. The path must be consistent across all nodes in the cluster.
The script creates a small partition on the device to be used for SBD.

5. Configure a virtual IP address for cluster administration with Hawk2:

- a. Confirm with `y` that you want to configure a virtual IP address.
- b. Enter an unused IP address that you want to use as administration IP for Hawk2:
`192.168.2.1`

Instead of logging in to an individual cluster node with Hawk2, you can connect to the virtual IP address.

6. Choose whether to configure QDevice and QNetd. For the minimal setup described in this document, decline with `n` for now. You can set up QDevice and QNetd later, as described in *Book "Administration Guide", Chapter 13 "QDevice and QNetd"*.

Finally, the script will start the Pacemaker service to bring the cluster online and enable Hawk2. The URL to use for Hawk2 is displayed on the screen.

You now have a running one-node cluster. To view its status, proceed as follows:

PROCEDURE 4: LOGGING IN TO THE HAWK2 WEB INTERFACE

1. On any machine, start a Web browser and make sure that JavaScript and cookies are enabled.

2. As URL, enter the virtual IP address that you configured with the bootstrap script:

```
https://192.168.2.1:7630/
```



Note: Certificate Warning

If a certificate warning appears when you try to access the URL for the first time, a self-signed certificate is in use. Self-signed certificates are not considered trustworthy by default.

Ask your cluster operator for the certificate details to verify the certificate.

To proceed anyway, you can add an exception in the browser to bypass the warning.

3. On the Hawk2 login screen, enter the *Username* and *Password* of the user that was created by the bootstrap script (user `hacluster`, password `linux`).



Important: Secure Password

Replace the default password with a secure one as soon as possible:

```
# passwd hacluster
```

4. Click *Log In*. The Hawk2 Web interface shows the Status screen by default:

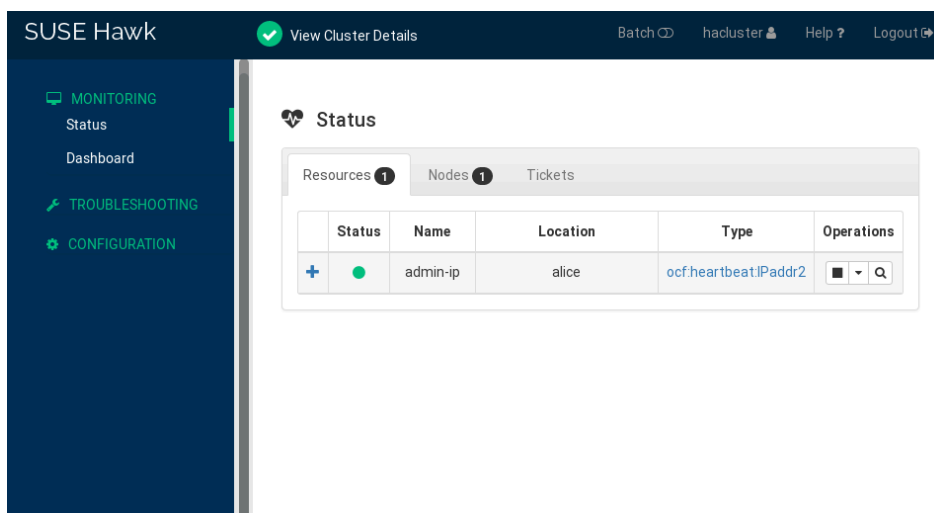


FIGURE 1: STATUS OF THE ONE-NODE CLUSTER IN HAWK2

7 Adding the Second Node

Add a second node to the cluster with the `ha-cluster-join` bootstrap script. The script only needs access to an existing cluster node and will complete the basic setup on the current machine automatically.

For more information, see the `ha-cluster-join` man page.

PROCEDURE 5: ADDING THE SECOND NODE (bob) WITH `ha-cluster-join`

1. Log in as `root` to the physical or virtual machine you want to add to the cluster.
2. Start the bootstrap script:

```
# ha-cluster-join
```

If NTP has not been configured to start at boot time, a message appears. The script also checks for a hardware watchdog device. You are warned if none is present.

3. If you decide to continue anyway, you will be prompted for the IP address of an existing node. Enter the IP address of the first node (`alice`, `192.168.1.1`).
4. If you have not already configured passwordless SSH access between both machines, you will be prompted for the `root` password of the existing node.
After logging in to the specified node, the script will copy the Corosync configuration, configure SSH and Csync2, bring the current machine online as new cluster node, and start the service needed for Hawk2.

Check the cluster status in Hawk2. Under *Status* > *Nodes* you should see two nodes with a green status:




FIGURE 2: STATUS OF THE TWO-NODE CLUSTER

8 Testing the Cluster

The following tests can help you identify issues with the cluster setup. However, a realistic test involves specific use cases and scenarios. Before using the cluster in a production environment, test it thoroughly according to your use cases.

- The command **`sbd -d DEVICE_NAME list`** lists all of the nodes that are visible to SBD. For the setup described in this document, the output should show both alice and bob.
- [Section 8.1, “Testing Resource Failover”](#) is a simple test to check if the cluster moves the virtual IP address to the other node if the node that currently runs the resource is set to standby.
- [Section 8.2, “Testing with the `crm cluster crash_test` command”](#) simulates cluster failures and reports the results.

8.1 Testing Resource Failover

As a quick test, the following procedure checks on resource failovers:

PROCEDURE 6: TESTING RESOURCE FAILOVER

1. Open a terminal and ping `192.168.2.1`, your virtual IP address:

```
# ping 192.168.2.1
```

2. Log in to Hawk2.
3. Under *Status* > *Resources*, check which node the virtual IP address (resource `admin_addr`) is running on. This procedure assumes the resource is running on `alice`.
4. Put `alice` into *Standby* mode:




FIGURE 3: NODE `alice` IN STANDBY MODE

5. Click *Status* > *Resources*. The resource `admin_addr` has been migrated to `bob`.

During the migration, you should see an uninterrupted flow of pings to the virtual IP address. This shows that the cluster setup and the floating IP work correctly. Cancel the `ping` command with `Ctrl-C`.

8.2 Testing with the `crm cluster crash_test` command

The command `crm cluster crash_test` triggers cluster failures to find problems. Before you use your cluster in production, it is recommended to use this command to make sure everything works as expected.

The command supports the following checks:

`--split-brain-iptables`

Simulates a split brain scenario by blocking the Corosync port. Checks whether one node can be fenced as expected.

--kill-sbd / --kill-corosync / --kill-pacemakerd

Kills the daemons for SBD, Corosync, and Pacemaker. After running one of these tests, you can find a report in the directory /var/lib/crmsh/crash_test/. The report includes a test case description, action logging, and an explanation of possible results.

--fence-node *NODE*

Fences a specific node passed from the command line.

For more information, see **crm cluster crash_test --help**.

EXAMPLE 1: TESTING THE CLUSTER: NODE FENCING

```
# crm_mon -l
Stack: corosync
Current DC: alice (version ...) - partition with quorum
Last updated: Fri Mar 03 14:40:21 2020
Last change: Fri Mar 03 14:35:07 2020 by root via cibadmin on alice

2 nodes configured
1 resource configured

Online: [ alice bob ]
Active resources:

    stonith-sbd    (stonith:external/sbd): Started alice

# crm cluster crash_test --fence-node bob

=====
Testcase:          Fence node bob
Fence action:      reboot
Fence timeout:     60

!!! WARNING WARNING WARNING !!!
THIS CASE MAY LEAD TO NODE BE FENCED.
TYPE Yes TO CONTINUE, OTHER INPUTS WILL CANCEL THIS CASE [Yes/No](No): Yes
INFO: Trying to fence node "bob"
INFO: Waiting 60s for node "bob" reboot...
INFO: Node "bob" will be fenced by "alice"!
INFO: Node "bob" was successfully fenced by "alice"
```

To watch bob change status during the test, log in to Hawk2 and navigate to *Status > Nodes*.

9 Next Steps

The bootstrap scripts provide a quick way to set up a basic High Availability cluster that can be used for testing purposes. However, to expand this cluster into a functioning High Availability cluster that can be used in production environments, more steps are recommended.

RECOMMENDED STEPS TO COMPLETE THE HIGH AVAILABILITY CLUSTER SETUP

Adding more nodes

Add more nodes to the cluster using one of the following methods:

- For individual nodes, use the `crm cluster join` script as described in [Section 7, “Adding the Second Node”](#).
- For mass installation of multiple nodes, use AutoYaST as described in *Book “Administration Guide”, Chapter 3 “Installing the High Availability Extension”, Section 3.2 “Mass Installation and Deployment with AutoYaST”*.

A regular cluster can contain up to 32 nodes. With the `pacemaker_remote` service, High Availability clusters can be extended to include additional nodes beyond this limit. See *Article “Pacemaker Remote Quick Start”* for more details.

Configuring QDevice

If the cluster has an even number of nodes, configure QDevice and QNetd to participate in quorum decisions. QDevice provides a configurable number of votes, allowing a cluster to sustain more node failures than the standard quorum rules allow. For details, see *Book “Administration Guide”, Chapter 13 “QDevice and QNetd”*.

Enabling a hardware watchdog

Before using the cluster in a production environment, replace the `softdog` module with the hardware module that best fits your hardware. For details, see *Book “Administration Guide”, Chapter 12 “Storage Protection and SBD”, Section 12.6 “Setting Up the Watchdog”*.

10 For More Information

More documentation for this product is available at <https://documentation.suse.com/sle-ha/15-GA/>. For further configuration and administration tasks, see the comprehensive *Administration Guide* (<https://documentation.suse.com/sle-ha/15-GA/html/SLE-HA-all/book-sleha-guide.html>).

This appendix contains the GNU Free Documentation License version 1.2.

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. Secondly, 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 <http://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.