

Highly Available NFS Storage with DRBD and Pacemaker

This document describes how to set up highly available NFS storage in a two-node cluster, using the following components: DRBD* (Distributed Replicated Block Device), LVM (Logical Volume Manager), and Pacemaker as cluster resource manager.

Publication Date: December 12, 2024

Contents

- 1 Usage scenario 2
- 2 Preparing a two-node cluster 2
- 3 Creating LVM devices 3
- 4 Creating DRBD devices 4
- 5 Creating cluster resources 7
- 6 Using the NFS service 12
- 7 Adding more NFS shares to the cluster 13
- 8 For more information 15
- A GNU licenses 16

1 Usage scenario

This document helps you set up a highly available NFS server. The cluster used for the highly available NFS storage has the following properties:


- Two nodes: `alice` (IP: `192.168.1.1`) and `bob` (IP: `192.168.1.2`), connected to each other via network.
- Two floating, virtual IP addresses (`192.168.1.10` and `192.168.1.11`), allowing clients to connect to a service no matter which physical node it is running on. One IP address is used for cluster administration with Hawk2, and the other IP address is used exclusively for the NFS exports.
- SBD used as a STONITH fencing device to avoid split-brain scenarios. STONITH is mandatory for the HA cluster.
- Failover of resources from one node to the other if the active host breaks down (*active/passive* setup).
- Local storage on each node. The data is synchronized between the nodes using DRBD on top of LVM.
- A file system exported through NFS and a separate file system used to track the NFS client states.

After installing and setting up the basic two-node cluster, and extending it with storage and cluster resources for NFS, you will have a highly available NFS storage server.

2 Preparing a two-node cluster

Before you can set up highly available NFS storage, you must prepare a High Availability cluster:

PROCEDURE 1: PREPARING A TWO-NODE CLUSTER FOR NFS STORAGE

1. Install and set up a basic two-node cluster as described in [Installation and Setup Quick Start](https://documentation.suse.com/sle-ha/15-SP5/html/SLE-HA-all/article-installation.html) (<https://documentation.suse.com/sle-ha/15-SP5/html/SLE-HA-all/article-installation.html>) .
2. On *both* nodes, install the package `nfs-kernel-server`:

```
# zypper install nfs-kernel-server
```

3 Creating LVM devices

LVM (Logical Volume Manager) enables flexible distribution of storage space across several file systems.

Use `crm cluster run` to run these commands on both nodes at once.

PROCEDURE 2: CREATING LVM DEVICES FOR DRBD

1. Create an LVM physical volume, replacing `/dev/disk/by-id/DEVICE_ID` with your corresponding device for LVM:

```
# crm cluster run "pvcreate /dev/disk/by-id/DEVICE_ID"
```

2. Create an LVM volume group `nfs` that includes this physical volume:

```
# crm cluster run "vgcreate nfs /dev/disk/by-id/DEVICE_ID"
```

3. Create a logical volume named `share` in the volume group `nfs`:

```
# crm cluster run "lvcreate -n share -L 20G nfs"
```

This volume is for the NFS exports.

4. Create a logical volume named `state` in the volume group `nfs`:

```
# crm cluster run "lvcreate -n state -L 8G nfs"
```

This volume is for the NFS client states. The 8 GB volume size used in this example should support several thousand concurrent NFS clients.

5. Activate the volume group:

```
# crm cluster run "vgchange -ay nfs"
```

You should now see the following devices on the system: `/dev/nfs/share` and `/dev/nfs/state`.

4 Creating DRBD devices

This section describes how to set up DRBD devices on top of LVM. Using LVM as a back-end of DRBD has the following benefits:

- Easier setup than with LVM on top of DRBD.
- Easier administration in case the LVM disks need to be resized or more disks are added to the volume group.

The following procedures result in two DRBD devices: one device for the NFS exports, and a second device to track the NFS client states.

4.1 Creating the DRBD configuration

DRBD configuration files are kept in the `/etc/drbd.d/` directory and must end with a `.res` extension. In this procedure, the configuration file is named `/etc/drbd.d/nfs.res`.

PROCEDURE 3: CREATING A DRBD CONFIGURATION

1. Create the file `/etc/drbd.d/nfs.res` with the following contents:

```
resource nfs {  
  volume 0 { ❶  
    device      /dev/drbd0; ❷  
    disk        /dev/nfs/state; ❸  
    meta-disk    internal; ❹  
  }  
  volume 1 {  
    device      /dev/drbd1;  
    disk        /dev/nfs/share;  
    meta-disk    internal;  
  }  
  
  net {  
    protocol C; ❺  
    fencing resource-and-stonith; ❻  
  }  
  
  handlers { ❼  
    fence-peer "/usr/lib/drbd/crm-fence-peer.9.sh";  
    after-resync-target "/usr/lib/drbd/crm-unfence-peer.9.sh";  
  }  
}
```

```

connection-mesh { ❸
    hosts    alice bob;
}
on alice { ❹
    address  192.168.1.1:7790;
    node-id  0;
}
on bob {
    address  192.168.1.2:7790;
    node-id  1;
}
}

```

- ❶ The volume number for each DRBD device you want to create.
- ❷ The DRBD device that applications will access.
- ❸ The lower-level block device used by DRBD to store the actual data. This is the LVM device that was created in [Section 3, “Creating LVM devices”](#).
- ❹ Where the metadata is stored. Using `internal`, the metadata is stored together with the user data on the same device. See the man page for further information.
- ❺ The protocol to use for this connection. Protocol `C` is the default option. It provides better data availability and does not consider a write to be complete until it has reached all local and remote disks.
- ❻ Specifies the fencing policy `resource-and-stonith` at the DRBD level. This policy immediately suspends active I/O operations until STONITH completes.
- ❼ Enables resource-level fencing to prevent Pacemaker from starting a service with outdated data. If the DRBD replication link becomes disconnected, the `crm-fence-peer.9.sh` script stops the DRBD resource from being promoted to another node until the replication link becomes connected again and DRBD completes its synchronization process.
- ❽ Defines all nodes of a mesh. The `hosts` parameter contains all host names that share the same DRBD setup.
- ❾ Contains the IP address and a unique identifier for each node.

2. Open `/etc/csync2/csync2.cfg` and check whether the following two lines exist:

```

include /etc/drbd.conf;
include /etc/drbd.d;

```

If not, add them to the file.

3. Copy the file to the other nodes:

```
# csync2 -xv
```

For information about Csync2, see Book “Administration Guide”, Chapter 4 “Using the YaST cluster module”, Section 4.7 “Transferring the configuration to all nodes”.

4.2 Activating the DRBD devices

After preparing the DRBD configuration, activate the devices:

PROCEDURE 4: ACTIVATING DRBD DEVICES

1. If you use a firewall in the cluster, open port 7790 in the firewall configuration.
2. Initialize the metadata storage:

```
# crm cluster run "drbdadm create-md nfs"
```

3. Create the DRBD devices:

```
# crm cluster run "drbdadm up nfs"
```

4. The devices do not have data yet, so you can run these commands to skip the initial synchronization:

```
# drbdadm new-current-uuid --clear-bitmap nfs/0  
# drbdadm new-current-uuid --clear-bitmap nfs/1
```

5. Make alice primary:

```
# drbdadm primary --force nfs
```

6. Check the DRBD status of nfs:

```
# drbdadm status nfs
```

This returns the following message:

```
nfs role:Primary  
  volume:0 disk:UpToDate  
  volume:1 disk:UpToDate  
bob  role:Secondary
```

```
volume:0 peer-disk:UpToDate
volume:1 peer-disk:UpToDate
```

You can access the DRBD resources on the block devices `/dev/drbd0` and `/dev/drbd1`.

4.3 Creating the file systems

After activating the DRBD devices, create file systems on them:

PROCEDURE 5: CREATING FILE SYSTEMS FOR DRBD

1. Create an `ext4` file system on `/dev/drbd0`:

```
# mkfs.ext4 /dev/drbd0
```

2. Create an `ext4` file system on `/dev/drbd1`:

```
# mkfs.ext4 /dev/drbd1
```

5 Creating cluster resources

The following procedures describe how to configure the resources required for a highly available NFS cluster.

OVERVIEW OF CLUSTER RESOURCES

DRBD primitive and promotable clone resources

These resources are used to replicate data. The promotable clone resource is switched to and from the primary and secondary roles as deemed necessary by the cluster resource manager.

File system resources

These resources manage the file system that will be exported, and the file system that will track NFS client states.

NFS kernel server resource

This resource manages the NFS server daemon.

NFS exports

This resource is used to export the directory `/srv/nfs/share` to clients.

Virtual IP address

The initial installation creates an administrative virtual IP address for Hawk2. Create another virtual IP address exclusively for NFS exports. This makes it easier to apply security restrictions later.

EXAMPLE NFS SCENARIO

- The following configuration examples assume that 192.168.1.11 is the virtual IP address to use for an NFS server which serves clients in the 192.168.1.x/24 subnet.
- The service exports data served from /srv/nfs/share.
- Into this export directory, the cluster mounts an ext4 file system from the DRBD device /dev/drbd1. This DRBD device sits on top of an LVM logical volume named /dev/nfs/share.
- The DRBD device /dev/drbd0 is used to share the NFS client states from /var/lib/nfs. This DRBD device sits on top of an LVM logical volume named /dev/nfs/state.

5.1 Creating DRBD primitive and promotable clone resources

Create a cluster resource to manage the DRBD devices, and a promotable clone to allow this resource to run on both nodes:

PROCEDURE 6: CREATING A DRBD RESOURCE FOR NFS

1. Start the crm interactive shell:

```
# crm configure
```

2. Create a primitive for the DRBD configuration nfs:

```
crm(live)configure# primitive drbd-nfs ocf:linbit:drbd \  
  params drbd_resource="nfs" \  
  op monitor interval=15 role=Promoted \  
  op monitor interval=30 role=Unpromoted
```

3. Create a promotable clone for the drbd-nfs primitive:

```
crm(live)configure# clone cl-drbd-nfs drbd-nfs \  
  meta promotable="true" promoted-max="1" promoted-node-max="1" \  
  clone-max="2" clone-node-max="1" notify="true" interleave=true
```


4. Commit this configuration:

```
crm(live)configure# commit
```

Pacemaker activates the DRBD resources on both nodes and promotes them to the primary role on one of the nodes. Check the state of the cluster with the `crm status` command, or run `drbdadm status`.

5.2 Creating file system resources

Create cluster resources to manage the file systems for export and state tracking:

PROCEDURE 7: CREATING FILE SYSTEM RESOURCES FOR NFS

1. Create a primitive for the NFS client states on `/dev/drbd0`:

```
crm(live)configure# primitive fs-nfs-state Filesystem \  
params device=/dev/drbd0 directory=/var/lib/nfs fstype=ext4
```

2. Create a primitive for the file system to be exported on `/dev/drbd1`:

```
crm(live)configure# primitive fs-nfs-share Filesystem \  
params device=/dev/drbd1 directory=/srv/nfs/share fstype=ext4
```

Do not commit this configuration until after you add the colocation and order constraints.

3. Add both of these resources to a resource group named `g-nfs`:

```
crm(live)configure# group g-nfs fs-nfs-state fs-nfs-share
```

Resources start in the order they are added to the group and stop in reverse order.

4. Add a colocation constraint to make sure that the resource group always starts on the node where the DRBD promotable clone is in the primary role:

```
crm(live)configure# colocation col-nfs-on-drbd inf: g-nfs cl-drbd-nfs:Promoted
```

5. Add an order constraint to make sure the DRBD promotable clone always starts before the resource group:

```
crm(live)configure# order o-drbd-before-nfs Mandatory: cl-drbd-nfs:promote g-  
nfs:start
```

6. Commit this configuration:

```
crm(live)configure# commit
```

Pacemaker mounts `/dev/drbd0` to `/var/lib/nfs`, and `/dev/drbd1` to `srv/nfs/share`. Confirm this with `mount`, or by looking at `/proc/mounts`.

5.3 Creating an NFS kernel server resource

Create a cluster resource to manage the NFS server daemon:

PROCEDURE 8: CREATING AN NFS KERNEL SERVER RESOURCE

1. Create a primitive to manage the NFS server daemon:

```
crm(live)configure# primitive nfsserver nfsserver \  
    params nfs_server_scope=SUSE nfs_shared_infodir="/var/lib/nfs"
```

The `nfs_server_scope` must be the same on all nodes in the cluster that run the NFS server, but this is not set by default. All clusters using SUSE software can use the same scope, so we recommend setting the value to `SUSE`.



Warning: Low lease time can cause loss of file state

NFS clients regularly renew their state with the NFS server. If the lease time is too low, system or network delays can cause the timer to expire before the renewal is complete. This can lead to I/O errors and loss of file state.

`NFSV4LEASETIME` is set on the NFS server in the file `/etc/sysconfig/nfs`. The default is 90 seconds. If lowering the lease time is necessary, we recommend a value of 60 or higher. We strongly discourage values lower than 30.

2. Append this resource to the existing `g-nfs` resource group:

```
crm(live)configure# modgroup g-nfs add nfsserver
```

3. Commit this configuration:

```
crm(live)configure# commit
```

5.4 Creating an NFS export resource

Create a cluster resource to manage the NFS exports:

PROCEDURE 9: CREATING AN NFS EXPORT RESOURCE

1. Create a primitive for the NFS exports:

```
crm(live)configure# primitive exportfs-nfs exportfs \  
  params directory="/srv/nfs/share" \  
  options="rw,mountpoint" clientspec="192.168.1.0/24" fsid=101 \① \  
  op monitor interval=30s timeout=90s②
```

- ① The `fsid` must be unique for each NFS export resource.
- ② The value of `op monitor timeout` must be higher than the value of `stonith-timeout`. To find the `stonith-timeout` value, run `crm configure show` and look under the `property` section.



Important: Do not set `wait_for_leasetime_on_stop=true`

Setting this option to `true` in a highly available NFS setup can cause unnecessary delays and loss of locks.

The default value for `wait_for_leasetime_on_stop` is `false`. There is no need to set it to `true` when `/var/lib/nfs` and `nfsserver` are configured as described in this guide.

2. Append this resource to the existing `g-nfs` resource group:

```
crm(live)configure# modgroup g-nfs add exportfs-nfs
```

3. Commit this configuration:

```
crm(live)configure# commit
```

4. Confirm that the NFS exports are set up properly:

```
# exportfs -v  
/srv/nfs/share IP_ADDRESS_OF_CLIENT(OPTIONS)
```

5.5 Creating a virtual IP address for NFS exports

Create a cluster resource to manage the virtual IP address for the NFS exports:

PROCEDURE 10: CREATING A VIRTUAL IP ADDRESS FOR NFS EXPORTS

1. Create a primitive for the virtual IP address:

```
crm(live)configure# primitive vip-nfs IPAddr2 params ip=192.168.1.11
```

2. Append this resource to the existing `g-nfs` resource group:

```
crm(live)configure# modgroup g-nfs add vip-nfs
```

3. Commit this configuration:

```
crm(live)configure# commit
```

4. Leave the `crm` interactive shell:

```
crm(live)configure# quit
```

5. Check the status of the cluster. The resources in the `g-nfs` group should appear in the following order:

```
# crm status
[...]
Full List of Resources
[...]
* Resource Group: g-nfs:
  * fs-nfs-state      (ocf:heartbeat:Filesystem): Started alice
  * fs-nfs-share      (ocf:heartbeat:Filesystem): Started alice
  * nfsserver         (ocf:heartbeat:nfsserver): Started alice
  * exportfs-nfs      (ocf:heartbeat:exportfs): Started alice
  * vip-nfs           (ocf:heartbeat:IPAddr2): Started alice
```

6 Using the NFS service

This section outlines how to use the highly available NFS service from an NFS client.

To connect to the NFS service, make sure to use the *virtual IP address* to connect to the cluster rather than a physical IP configured on one of the cluster nodes' network interfaces. For compatibility reasons, use the *full* path of the NFS export on the server.

The command to mount the NFS export looks like this:

```
# mount 192.168.1.11:/srv/nfs/share /home/share
```

If you need to configure other mount options, such as a specific transport protocol (proto), maximum read and write request sizes (rsize and wsizes), or a specific NFS version (vers), use the -o option. For example:

```
# mount -o proto=tcp,rsize=32768,wsizes=32768,vers=3 \
192.168.1.11:/srv/nfs/share /home/share
```

For further NFS mount options, see the nfs man page.



Note: Loopback mounts

Loopback mounts are only supported for NFS version 3, *not* NFS version 4. For more information, see <https://www.suse.com/support/kb/doc/?id=000018709>.

7 Adding more NFS shares to the cluster

If you need to increase the available storage, you can add more NFS shares to the cluster.

In this example, a new DRBD device named /dev/drbd2 sits on top of an LVM logical volume named /dev/nfs/share2.

PROCEDURE 11: ADDING MORE NFS SHARES TO THE CLUSTER

1. Create an LVM logical volume for the new share:

```
# crm cluster run "lvcreate -n share2 -L 20G nfs"
```

2. Update the file /etc/drbd.d/nfs.res to add the new volume under the existing volumes:

```
volume 2 {
    device          /dev/drbd2;
    disk            /dev/nfs/share2;
    meta-disk       internal;
}
```

3. Copy the updated file to the other nodes:

```
# csync2 -xv
```

4. Initialize the metadata storage for the new volume:

```
# crm cluster run "drbdadm create-md nfs/2 --force"
```

5. Update the `nfs` configuration to create the new device:

```
# crm cluster run "drbdadm adjust nfs"
```

6. Skip the initial synchronization for the new device:

```
# drbdadm new-current-uuid --clear-bitmap nfs/2
```

7. The NFS cluster resources might have moved to another node since they were created. Check the DRBD status with `drbdadm status nfs`, and make a note of which node is in the `Primary` role.

8. On the node that is in the `Primary` role, create an `ext4` file system on `/dev/drbd2`:

```
# mkfs.ext4 /dev/drbd2
```

9. Start the `crm` interactive shell:

```
# crm configure
```

10. Create a primitive for the file system to be exported on `/dev/drbd2`:

```
crm(live)configure# primitive fs-nfs-share2 Filesystem \  
  params device="/dev/drbd2" directory="/srv/nfs/share2" fstype=ext4
```

11. Add the new file system resource to the `g-nfs` group *before* the `nfsserver` resource:

```
crm(live)configure# modgroup g-nfs add fs-nfs-share2 before nfsserver
```

12. Create a primitive for NFS exports from the new share:

```
crm(live)configure# primitive exportfs-nfs2 exportfs \  
  params directory="/srv/nfs/share2" \  
  options="rw,mountpoint" clientspec="192.168.1.0/24" fsid=102 \  
  op monitor interval=30s timeout=90s
```

13. Add the new NFS export resource to the `g-nfs` group *before* the `vip-nfs` resource:

```
crm(live)configure# modgroup g-nfs add exportfs-nfs2 before vip-nfs
```

14. Commit this configuration:

```
crm(live)configure# commit
```

15. Leave the `crm` interactive shell:

```
crm(live)configure# quit
```

16. Check the status of the cluster. The resources in the `g-nfs` group should appear in the following order:

```
# crm status
[...]
Full List of Resources
[...]
* Resource Group: g-nfs:
  * fs-nfs-state      (ocf:heartbeat:Filesystem): Started alice
  * fs-nfs-share      (ocf:heartbeat:Filesystem): Started alice
  * fs-nfs-share2     (ocf:heartbeat:Filesystem): Started alice
  * nfsserver         (ocf:heartbeat:nfsserver):  Started alice
  * exportfs-nfs      (ocf:heartbeat:exportfs):   Started alice
  * exportfs-nfs2     (ocf:heartbeat:exportfs):   Started alice
  * vip-nfs           (ocf:heartbeat:IPAddr2):    Started alice
```

17. Confirm that the NFS exports are set up properly:

```
# exportfs -v
/srv/nfs/share      IP_ADDRESS_OF_CLIENT(OPTIONS)
/srv/nfs/share2     IP_ADDRESS_OF_CLIENT(OPTIONS)
```

8 For more information

- For more details about the steps in this guide, see <https://www.suse.com/support/kb/doc?id=000020396>.
- For more information about NFS and LVM, see [Storage Administration Guide for SUSE Linux Enterprise Server \(https://documentation.suse.com/sles/html/SLES-all/book-storage.html\)](https://documentation.suse.com/sles/html/SLES-all/book-storage.html).
- For more information about DRBD, see *Book "Administration Guide", Chapter 23 "DRBD"*.
- For more information about cluster resources, see *Book "Administration Guide", Chapter 6 "Configuring cluster resources"*.

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 <https://www.gnu.org/copyleft/ 3>.

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.