

Data Replication across Geo Clusters via DRBD

Included with SUSE Linux Enterprise High Availability
Extension

SUSE Linux Enterprise High Availability Extension 12 SP1,
SP2

Matt Kereczman, Cluster Engineer, LINBIT

Philipp Marek, Senior Software Developer, LINBIT

Kristoffer Grönlund, Architect High Availability, SUSE

This technical guide describes the setup of a geo cluster using DRBD as delivered with the SUSE Linux Enterprise High Availability Extension.

Publication Date: November 7, 2016

Contents


- 1 Background 3
- 2 Introduction 4
- 3 Installation Requirements 4
- 4 Cluster Architecture Overview 5
- 5 Example Configurations for a Single Service 6
- 6 Interoperability with Booth 11
- 7 Testing Failovers and Recovery 13
- 8 Further Documentation 14
- 9 Legal Notice 14
- 10 GNU Free Documentation License 16

1 Background

The following sections provide you with important background information about DRBD and the SUSE Linux Enterprise High Availability Extension.

1.1 About DRBD

DRBD is a Linux-kernel block-level replication facility that is widely used as a shared-nothing cluster building block. It is included in vanilla kernels since 2.6.33, and most distributions ship the necessary userspace utilities. Furthermore, many distributions have newer DRBD versions than the one included in the kernel in extra packages.

DRBD can replicate across multiple network protocols, and in (currently) three modes, from synchronous for local HA clusters, to asynchronous for pushing data to a disaster recovery site. DRBD is developed and supported worldwide by LINBIT <http://www.linbit.com> . That includes most distributions and architectures, with Service Level Agreement (SLA) levels up to 24/7 e-mail and phone availability.

1.2 About SUSE Linux Enterprise High Availability Extension

SUSE Linux Enterprise High Availability Extension is an integrated suite of open source clustering technologies that enables you to implement highly available physical and virtual Linux clusters, and to eliminate single points of failure. It ensures the availability and manageability of critical networked resources including data, applications, and services. Thus, it helps you maintain business continuity, protect data integrity, and reduce unplanned downtime for your mission-critical Linux workloads.

SUSE Linux Enterprise High Availability Extension ships with essential monitoring, messaging, and cluster resource management functionality (supporting failover, failback, and migration (load balancing) of individually managed cluster resources).

SUSE Linux Enterprise High Availability Extension includes DRBD.

2 Introduction

This technical guide describes a setup consisting of two highly available Pacemaker clusters in two sites, with a routed IPv4 or IPv6 connection in between. The connection can range from a few Mbit/sec up to 10 GBit/sec, depending on the IO load imposed on the cluster nodes.

Various services can be distributed across the clusters. Because of latency between the data centers you cannot drive a cluster file system across them. But simply replicating the data to have a quick failover in case one site goes down is possible.

Failover between the two sites is managed using the *Booth Cluster Ticket Manager*, which is included in the Geo Clustering for SUSE Linux Enterprise High Availability Extension and the LINBIT HA Cluster Stack. In addition to the two cluster sites, a third site is needed to run a booth Arbitrator. Arbitrators are single machines that run a booth instance in a special mode. As all booth instances communicate with each other, arbitrators help to make more reliable decisions about granting or revoking tickets.

Lots of details will be skimmed. For example, for this technical guide it does not matter whether the application driving IO is SAP, an NFS server, a PostgreSQL instance, VMs via KVM, Apache, MySQL, or whatever else you may have in use.

3 Installation Requirements

You need to have the Geo Clustering for SUSE Linux Enterprise High Availability Extension installed on both pairs of nodes, and on the Arbitrator node. For a detailed description on how to configure the geo cluster nodes, see the Geo Clustering Quick Start document included in the Geo Clustering for SUSE Linux Enterprise High Availability Extension documentation at https://www.suse.com/documentation/sle-ha-geo-12/art_ha_geo_quick/data/art_ha_geo_quick.html.

It is good practice and recommended to use LVM as storage layer below DRBD. This allows for quick and easy creation and online enlarging of DRBD devices.

You also need to install all the software including dependencies for the services you want to run on all four nodes. In case you replicate Virtual Machines (VMs), these are typically self-contained, so you will just need the KVM environment *libvirt* with *virsh* on the cluster nodes.

4 Cluster Architecture Overview

The following image depicts a two-site, four-node cluster with arrows showing the direction of replication, and blocks in orange showing that a DRBD device is Primary. There are two nodes in the local site, *alice* and *bob*, and two nodes in the remote site, *charlie* and *daisy*.

Each site will have its own DRBD resource stack and service IP to allow independent migration across the cluster nodes.

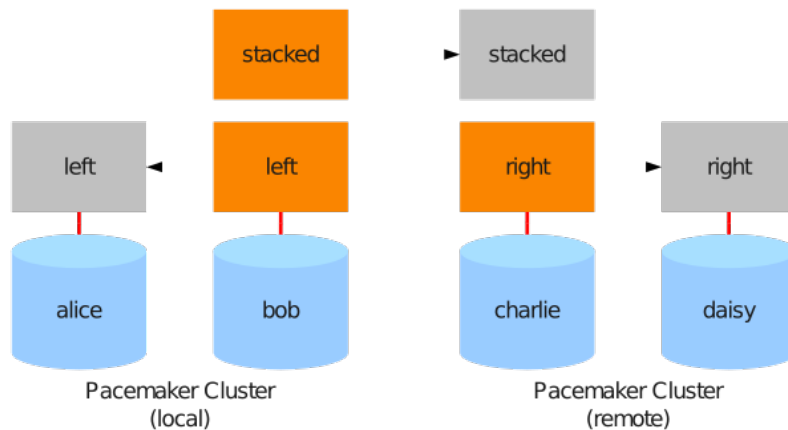


FIGURE 1: STACKED FOUR-NODE CLUSTER

The most important details of the cluster architecture are the following:

- The local high availability setup: DRBD running in Protocol C, using IP addresses in a LAN respective a cross-over connection.
- Just after activating the lower DRBD devices on one node, the dedicated service IP address is started.
This is not only used for the service as such, but also as a fixed point that can be accessed by the upper DRBD device (in Secondary state) for replication.
- The upper layer DRBD runs on one node per site, and is responsible for replicating the data to the other replication site. This connection runs in protocol A, and might have a DRBD Proxy setup in between.



Note: Data Compression

The DRBD Proxy buffers and optionally compresses data, from one or multiple DRBD resources. Compression helps if the site interconnect is slower than the summed average I/O rate of the resources. The ratio is about 1:4 for **zlib**, and up to 1:50 for **lzma** compression. See the “DRBD Proxy 3 Compression Considerations” whitepaper on https://links.linbit.com/tech_guides for more details.

- On the site that should actually run the service, the upper layer DRBD gets set as Primary, so that the file system therein can be mounted and used by the application.

5 Example Configurations for a Single Service

The example configurations below are using the following premises:

- Two sites, called **RZ1** and **RZ2**, with two private networks **192.168.201.x** and **192.168.202.x**, routed to the other site
- Four nodes, called **geo_rzN-M** in four combinations: **geo-rz1-a** to **geo-rz2-b**
- NFS is to be served; but there’s not much difference for other services
- Nodes are using LVM, VG name is **volgroup**
- The lower DRBD layer (for the HA-clusters) uses minor 0; minor 10 is used for DR replication

5.1 DRBD Configuration

The following snippets show a basic DRBD configuration. These are bare-bones; performance-tuning options are not included here.

All three snippets can be contained in a single resource file, for example in `/etc/drbd.d/nfs.res`. This is the recommended configuration, because keeping a single file synchronized across the four cluster nodes is easier. You can also consult **csync2** at <http://oss.linbit.com/csync2/>.

If you have used the Geo Clustering Quick Start Guide at https://www.suse.com/documentation/sle-ha-geo-12/art_ha_geo_quick/data/art_ha_geo_quick.html to perform the basic configuration of the cluster nodes, the DRBD configuration files are already included in the list of files to be synchronized.

To synchronize any changes to the configuration files across both cluster nodes, use the following command:

```
root # csync2 -xv /etc/drbd.d/
```

If you do not have `csync2`, or if you do not want to use it, you will need to copy the DRBD configuration files manually to all the other nodes.

5.1.1 DRBD on Site 1

To configure DRBD on Site 1, you should be aware of the following details:

- The resource-name has the site in it, so that the complete configuration can be kept in synchronization across both clusters without naming conflicts
- The nodes' **local** per-node IP addresses are used
- A *shared-secret* is used to avoid inadvertent wrong connections



Note: `uuid` Program

The `uuid` program is an easy way to get unique values.

```
resource nfs-lower-rz1 {
    disk          /dev/volgroup/lv-nfs;
    meta-disk     internal;
    device        /dev/drbd0;
    protocol      C;

    net {
        shared-secret "2a9702a6-8747-11e3-9ebb-782bcbd0c11c";
    }

    on geo-rz1-a {
        address      192.168.201.111:7900;
    }
    on geo-rz1-b {
```

```
        address      192.168.201.112:7900;
    }
}
```

5.1.2 DRBD on Site 2

Even if Site 2 is nearly identical to Site 1, you should notice the following differences:

- The resource name has changed
- The node names and IP addresses are different
- Another *shared-secret* has been generated
- The volume group and LV name can be kept in the *resource* section if identical on both nodes

```
resource nfs-lower-rz2 {
    disk          /dev/volgroup/lv-nfs;
    meta-disk     internal;
    device        /dev/drbd0;
    protocol      C;

    net {
        shared-secret "cd9d857d-72ef-4d10-a1de-6450d1797a2c";
    }

    on geo-rz2-a {
        address      192.168.202.111:7900;
    }
    on geo-rz2-b {
        address      192.168.202.112:7900;
    }
}
```

5.1.3 DRBD Connection across Sites

To configure a DRBD connection across sites, you should be aware of the following:

- The storage disk is the HA-cluster DRBD device */dev/drbd0*

- You could also use `/dev/drbd/by-res/nfs-lower-rzN/0`, but that would be site-specific, and so would need to be moved into the per-site configuration (stacked on top of `nfs-lower-rzN`)
- The DRBD device `drbd10` says to use minor number 10
- Protocol A, and a higher ping-timeout are needed because of the higher latency
- A different shared-secret is used
- Do not pass any host names, but tell DRBD to stack upon its lower device; that implies that this must be Primary
- To allow TCP/IP connections to the other site without knowing which cluster node has the lower DRBD device Primary, we are using a (the) service IP address

```
resource nfs-upper {
    disk          /dev/drbd0;
    meta-disk     internal;
    device        /dev/drbd10;
    protocol      A;

    net {
        shared-secret "e0fbd1fe-6b0b-47db-829a-2c4ba638bf1e";
        ping-timeout  20;
    }

    stacked-on-top-of nfs-lower-rz1 {
        address        192.168.201.151:7910;
    }
    stacked-on-top-of nfs-lower-rz2 {
        address        192.168.202.151:7910;
    }
}
```

Using a DRBD Proxy would involve inserting `proxy on ...` sections into `stacked-on-top-of` above, and a `proxy { ... }` section inside of `resource`. See LINBIT's DRBD Proxy guide at <http://links.linbit.com/DRBD-Proxy-Conf> for more details regarding configuring DRBD Proxy.

5.2 Pacemaker Resources (in crm-shell syntax)

For a more in-depth look at how to configure the NFS server, see the Highly Available NFS Storage with DRBD and Pacemaker document included in the SUSE Linux Enterprise High Availability Extension documentation. To configure the necessary resources, use the **crm** shell commands as outlined in the following chapters.

5.2.1 Basic Primitives

Setting up the basic primitives is fairly straightforward. You need a service IP, the file system, and the NFS server.



Note: exportfs

It is also possible to use the *exportfs* resource agents instead, and keep the NFS server running all the time. This is necessary if there are multiple NFS exports that must migrate independently.

```
crm configure
primitive p-ip-nfs IPAddr2 ip=192.168.202.151 iflabel=nfs nic=eth1 cidr_netmask=24
primitive p-nfs-fs Filesystem device=/dev/drbd/by-res/nfs/0 directory=/mnt/nfs
  fstype=ext4
primitive p-nfs-service systemd:nfs-server
```

5.2.2 DRBD Pacemaker Resources, Multi-State Resources

To configure the cluster resources for DRBD, it is possible to use the *drbd* cluster script. This script will create a base resource for DRBD, as well as a multi-state resource that ensures that DRBD only runs in Primary mode on a single node. Multi-state resources, previously called *Master-Slave* resources, allow the instances to be in one of two operating modes (called roles). The roles are called *master* and *slave*.

```
crm script run drbd id=drbd-nfs drbd_resource=nfs-upper
crm script run drbd id=drbd-nfs-lower drbd_resource=nfs-lower-rz2
```

5.2.3 Group and Basic Primitives

This is mostly what you would expect from the earlier picture: the multi-state equivalent of having a group consisting of *ms-drbd-nfs-lower:Master*, *p-ip-nfs*, and *ms-drbd-nfs:Master*.

```
crm configure

group g-nfs p-nfs-fs p-nfs-service

colocation co-nfs-ip-with-lower inf: p-ip-nfs:Started ms-drbd-nfs-lower:Master
colocation co-nfs-g-with-upper inf: g-nfs:Started ms-drbd-nfs:Master
colocation co-nfs-upper-with-ip inf: ms-drbd-nfs:Master p-ip-nfs:Started

order o-lower-drbd-before-ip-nfs inf: ms-drbd-nfs-lower:promote p-ip-nfs:start
order o-ip-nfs-before-drbd inf: p-ip-nfs:start ms-drbd-nfs:promote
order o-drbd-nfs-before-svc inf: ms-drbd-nfs:promote g-nfs:start
```

6 Interoperability with Booth

In the following sections, you can find the example configurations for the booth cluster ticket manager. Use *csync2* to synchronize the configurations for booth across all cluster nodes.

6.1 Booth Configuration

In the booth configuration, you need to specify the following components:

- A UDP port to use
- Three IP addresses



Note

You need one distinct service IP for each site, and a third one for the arbitrator. Separate addresses are preferred, so that booth can be managed independently.

- The Pacemaker ticket name (here *ticket-nfs*)

The configuration file might be stored as [/etc/booth/nfs.conf](#).

```
transport = udp
port      = "9929"

site      = "192.168.201.100"
site      = "192.168.202.100"
arbitrator = "192.168.203.100"

ticket = "ticket-nfs"
    expire = 600
    timeout = 5
    acquire-after = 60
```

6.2 Pacemaker Integration

Follow the recommendations below to integrate booth into Pacemaker:

- This example is again valid for site 2 (service IP address subnet).
- *booth* automatically uses the default directory and suffix if the configuration name does not specify any.
- As the DRBD replication should be running even if a site does not have the ticket, the correct *loss-policy* is *demote*. This will put DRBD into Secondary mode.
- As described in the **SUSE Multi-Site Cluster Documentation** at https://www.suse.com/documentation/sle_ha/book_sleha/data/sec_ha_geo_setup.html, there should be an order constraint that makes sure that Booth can fetch the ticket before trying to start the service.

```
crm configure

primitive p-booth ocf:pacemaker:booth-site config=nfs
primitive p-ip-booth IPAddr2 ip=192.168.202.100 iflabel=ha nic=eth1 cidr-netmask=24

group g-booth p-ip-booth p-booth

rsc-ticket nfs-req-ticket ticket-nfs: ms-drbd-nfs:Master loss-policy=demote

order o-booth-before-nfs inf: g-booth ms-drbd-nfs:promote
```

7 Testing Failovers and Recovery

After you have configured everything as described, you should check that your setup is working as expected.

The following command from any node running the booth daemon will tell you which tickets it knows about, which site currently holds the ticket, and when that ticket will expire if it is not renewed:

```
# booth list
ticket: ticket-nfs, leader: 192.168.201.100, expires: 2016-04-14 07:50:48
```

7.1 Failover

There are many ways to test that failovers will work as expected. An easy way to test is to sever network communications between sites using IPtables. You only need to cut communication on the Booth ports to see a ticket expire and services failover.

On the node currently running the booth daemon, issue the following IPtables commands to sever Booth's communications with its peers:

```
/usr/sbin/iptables -I INPUT -p udp --dport 9929 -j DROP
/usr/sbin/iptables -I OUTPUT -p udp --dport 9929 -j DROP
```

After the ticket's *expire* time has lapsed, you should see services begin to demote at the current Primary site. After the *expire* plus *acquire-after* time lapses, you should see services begin to start at the Secondary site.

Remove the IPtables rules you inserted before continuing:

```
# /usr/sbin/iptables -D INPUT -p udp --dport 9929 -j DROP
# /usr/sbin/iptables -D OUTPUT -p udp --dport 9929 -j DROP
```

7.2 Failback

To failback to the original Primary site you need to manually revoke the ticket from the site currently running services, and grant the ticket to the original Primary site.

You can issue the following commands from any of the nodes currently running the Booth daemon:

```
# booth revoke -s 192.168.202.100 ticket-nfs
  booth[30268]: 2016/04/14_08:21:05 info: revoke request sent, waiting for the
  result ...
  booth[30268]: 2016/04/14_08:21:09 info: revoke succeeded!
# booth grant -s 192.168.201.100 ticket-nfs
  booth[30269]: 2016/04/14_08:21:17 info: grant request sent, waiting for the
  result ...
  booth[30269]: 2016/04/14_08:21:28 info: grant succeeded
```

8 Further Documentation

SUSE Linux Enterprise High Availability Extension Guide: A comprehensive documentation about nearly every part of the Linux Cluster stack. See https://www.suse.com/documentation/sle_ha/singlehtml/book_sleha/book_sleha.html.

DRBD Project Page: Located at <https://www.drbd.org>, it provides lots of information, including a detailed Users' Guide, one of the most extensive project documentations in the Open Source world.

LINBIT Home Page: At <https://www.linbit.com>, you can find answers to all questions about paid support from the developers. An overview about supported platforms, SLAs, and price quotes is available at <https://www.linbit.com/en/p/services/support>.

SUSE Geo Clustering Documentation: This describes the general challenges for Geo-clustering, and typical solutions. The HTML version is hosted at https://www.suse.com/documentation/sle-ha-geo-12/singlehtml/art_ha_geo_quick/art_ha_geo_quick.html#sec.ha.geo.challenges.

9 Legal Notice

Copyright ©2006– 2016 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”.

SUSE, the SUSE logo and YaST are registered trademarks of SUSE LLC in the United States and other countries. For SUSE trademarks, see <http://www.suse.com/company/legal/> ↗.

DRBD and LINBIT are trademarks or registered trademarks of LINBIT in Austria, the United States, and other countries.

Linux is a registered trademark of Linus Torvalds. Other names or trademarks mentioned in this document may be trademarks or registered trademarks of their respective owners.

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.

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 noncommercially, 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.