

QEMU Copy-On-Write image file format specification

Analysis of the QCOW image format

By Joachim Metz <joachim.metz@gmail.com>

Summary

The QEMU Copy-On-Write (QCOW) image file format is used by the QEMU Open Source Process Emulator to store disk images.

This document is intended as a working document for the QCOW specification. Which should allow existing Open Source forensic tooling to be able to process this file type.

Document information

Author(s): Joachim Metz <joachim.metz@gmail.com>

Abstract: This document contains information about the QEMU Copy-On-Write file format.

Classification: Public

Keywords: QEMU Copy-On-Write file, qcow

License

Copyright (c) 2010-2013 Joachim Metz <joachim.metz@gmail.com>
Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 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".

Version

Version	Author	Date	Comments
0.0.1	J.B. Metz	December 2010 January 2011	Initial version
0.0.2	J.B. Metz	May 2012	Small update.
0.0.3	J.B. Metz	April 2013 May 2013	Additional information regarding QCOW version 1, compression and encryption.
0.0.4	J.B. Metz	May 2013	Additional information regarding QCOW version 2.
0.0.5	J.B. Metz	December 2013	Additional information regarding compressed QCOW version 2.

Table of Contents

1. Overview.....	1
1.1. Test version.....	1
2. Encryption.....	1
2.1. AES-CBC 128-bit.....	1
3. File header.....	2
3.1. File header – version 1.....	2
3.2. File header – version 2.....	3
3.3. Encryption methods.....	3
4. Cluster blocks.....	4
4.1. Level 1 table.....	4
4.1.1. Level 2 table reference – version 1.....	4
4.1.2. Level 2 table reference – version 2.....	4
4.2. Level 2 table.....	4
4.2.1. Cluster block reference – version 1.....	5
4.2.2. Cluster block reference – version 2.....	5
4.3. Reference count table.....	5
4.4. Retrieving a cluster data block.....	5
4.4.1. Uncompressed chunk data block.....	6
4.4.2. Compressed chunk data block.....	6
Compressed chunk data block – version 1.....	6
Compressed chunk data block – version 2.....	7
5. Snapshots.....	7
5.1. Snapshot header - version 2.....	7
6. Notes.....	8
6.1. Reference count table.....	8
Appendix A. References.....	9
Appendix B. GNU Free Documentation License.....	9

1. Overview

The QEMU Copy-On-Write (QCOW) image file format is used by the QEMU Open Source Process Emulator to store disk images (storage media).

The QCOW image file consists of:

- the file header
- the level 1 table (cluster aligned)
- the reference count table (cluster aligned)
- reference count blocks
- snapshot headers (8-byte aligned on cluster boundary)
- clusters containing:
 - level 2 tables
 - storage media data

The storage media data is stored in clusters. Each cluster is a multitude of 512 bytes. The level 1 (L1) table contains level 1 reference of level 2 (L2) tables. The level 2 tables contain level 2 references of the storage media clusters.

Characteristics	Description
Byte order	Big-endian in most cases, note that some values are in little-endian.
Date and time values	Number of seconds since Jan 1, 1970 00:00:00 UTC (POSIX epoch)
Character string	Narrow character strings assumed stored in UTF-8

1.1. Test version

The following version of programs were used to test the information within this document:

- QEMU

2. Encryption

The QCOW image format can encrypted the media data stored in the image format. Currently supported encryption methods are:

- AES-CBC 128-bit

If no encryption is used the encryption method in the file header is set to none (0).

It seems that compression and encryption at the same time is not supported, at least at the moment by qemu-img.

2.1. AES-CBC 128-bit

Both encryption and decryption use:

- AES-CBC with a 128-bits key decryption of sector data

The key is direct copy of the first 16 characters of a user provided (narrow character) password. If the password is smaller than 16 characters. The remaining key data is set to 0-byte values.

It is unclear which character sets are allowed and how characters outside the 7-bit ASCII set should be handled.

The initialization vector of the AES-CBC is using media data sector number (relative to the start of the disk) in little-endian format as the first 64 bits of the 128 bit initialization vector. The remaining initialization vector data is set to 0-byte values. The first sector number is 0 and the bytes per sector are 512.

3. File header

3.1. File header – version 1

The version 1 file header is 48 bytes of size and consist of:

offset	size	value	description
0	4	“QFI\xfb”	The signature (magic identifier) 0x51 0x46 0x49 0xfb
4	4	1	Version
8	8		Backing filename offset
16	4		Backing filename size
20	4		Modification date and time Contains a POSIX timestamp
24	8		Storage media size
32	1		Number of cluster block bits
33	1		Number of level 2 table bits
34	2		Unknown (empty values)
36	4		Encryption method
40	8		Level 1 table offset

The cluster block size is calculated as:

```
cluster block size = 1 << number of cluster block bits
```

The level table 2 size is calculated as:

```
level table 2 size = ( 1 << number of level 2 table bits ) * 8
```

The level 1 table size is calculated as:

```
level 1 table size = cluster block size
                    * ( 1 << number of level 2 table bits )

if( media size % level 1 table size != 0 )
{
    level 1 table size = ( media size / level 1 table size ) + 1
}
else
{
    level 1 table size = media size / level 1 table size
}
```

```
level 1 table size *= 8
```

The backing filename is set in snapshot image files and is normally stored after the file header.

3.2. File header – version 2

The version 2 file header is 72 bytes of size and consist of:

offset	size	value	description
0	4	“QFI\xfb”	The signature (magic identifier) 0x51 0x46 0x49 0xfb
4	4	2	Version
8	8		Backing filename offset
16	4		Backing filename size
20	4		Number of cluster block bits
24	8		Storage media size
32	4		Encryption method
36	4		Number of level 1 table references
40	8		Level 1 table offset
48	8		Reference count table offset
56	4		Reference count table clusters
60	4		Number of snapshots
64	8		Snapshots offset

The cluster block size is calculated as:

```
cluster block size = 1 << number of cluster block bits
```

The number of level 2 table bits is calculated as:

```
number of level 2 table bits = number of cluster block bits - 3
```

The level table 2 size is calculated as:

```
level table 2 size = ( 1 << number of level 2 table bits ) * 8
```

The level 1 table size is calculated as:

```
level 1 table size = number of level 1 table references * 8
```

The backing filename is set in snapshot image files and is normally stored after the file header.

3.3. Encryption methods

Value	Identifier	Description
0	QCOW_CRYPT_NONE	No encryption

Value	Identifier	Description
1	QCOW_CRYPT_AES	AES-CBC 128-bits encryption

4. Cluster blocks

4.1. Level 1 table

The level 1 table contains level 2 table references.

A reference value of 0 represents unused or unallocated and is considered as sparse or stored in a corresponding backing file.

4.1.1. Level 2 table reference – version 1

The level 2 table reference is 8-bytes of size and consists of:

offset	size	value	description
0.0	63 bits		Level 2 table offset Contains an offset value relative from the start of the file
7.7	1 bit		Is compressed flag (QCOW_OFLAG_COMPRESSED)

4.1.2. Level 2 table reference – version 2

The level 2 table reference is 8-bytes of size and consists of:

offset	size	value	description
0.0	62 bits		Level 2 table offset Contains an offset value relative from the start of the file
7.6	1 bit		Is compressed flag (QCOW_OFLAG_COMPRESSED)
7.7	1 bit		Is copied flag (QCOW_OFLAG_COPIED)

The is copied flag indicates that the reference count of the corresponding level 2 table is exactly one.

4.2. Level 2 table

The level 2 table contains cluster block references.

The level 2 table size is calculated as:

$$\text{level 2 table size} = (1 \ll \text{number of level 2 table bits}) * 8$$

A reference value of 0 represents unused or unallocated and is considered as sparse or stored in a corresponding backing file.

4.2.1. Cluster block reference – version 1

The cluster block reference is 8-bytes of size and consists of:

offset	size	value	description
0.0	63 bits		Cluster block offset The offset is relative to the start of the cluster block
7.7	1 bit		Is compressed flag (QCOW_OFLAG_COMPRESSED)

4.2.2. Cluster block reference – version 2

The cluster block reference is 8-bytes of size and consists of:

offset	size	value	description
0.0	62 bits		Cluster block offset The offset is relative to the start of the cluster block
7.6	1 bit		Is compressed flag (QCOW_OFLAG_COMPRESSED)
7.7	1 bit		Is copied flag (QCOW_OFLAG_COPIED)

The is copied flag indicates that the reference count of the corresponding cluster block is exactly one.

4.3. Reference count table

The cluster data blocks are referenced counted. For every cluster data block a 16-bit reference count is stored in the reference count table.

The reference count table is stored in cluster block sizes. The file header contains the number of blocks (or reference count table clusters).

TODO

4.4. Retrieving a cluster data block

To retrieve a cluster data block corresponding a certain storage media offset:

Determine the level 1 table index from the offset:

```
level 1 table index bit shift = number of cluster block bits
                               + number of level 2 table bits
```

Version 1

```
level 1 table index = ( offset & 0x7fffffffffffffffULL )
```

```
>> level 1 table index bit shift
```

Version 2

```
level 1 table index = ( offset & 0x3ffffffffffffffffULL )  
>> level 1 table index bit shift
```

Retrieve the level 2 table offset from the level 1 table. If the level table 2 offset is 0 and the image has a backing file the cluster data block is stored in the backing file otherwise the cluster block is considered sparse.

Read the corresponding level 2 table.

Determine the level 2 table index from the offset:

```
level 2 table index bit mask = ~( 0xffffffffffffffffULL  
                                  << number of level 2 table bits )
```

```
level 2 table index = ( offset >> number of cluster block bits )  
>> level 2 table index bit mask
```

Retrieve the cluster block offset from the level 2 table. If the cluster block offset is 0 and the image has a backing file the cluster data block is stored in the backing file otherwise the cluster block is considered sparse.

4.4.1. Uncompressed chunk data block

If the is compressed flag (QCOW_OFLAG_COMPRESSED) is not set:

```
cluster block bit mask = ~( 0xffffffffffffffffULL  
                             << number of cluster block bits )
```

```
cluster block data offset = ( offset & cluster block bit mask )  
                             + cluster block offset
```

Note that in version 2 the last cluster block in the file can be smaller than the cluster block size defined by the number of cluster block bits in the file header. This does not seem to be the case for version 1.

4.4.2. Compressed chunk data block

If the is compressed flag (QCOW_OFLAG_COMPRESSED) is set:

Compressed chunk data block – version 1

Version 1

```
compressed size bit shift = 63 - number of cluster block bits
```

```
compressed block size = ( ( cluster block offset & 0x7ffffffffffffffffULL )
```

```
>> compressed size bit shift
```

```
compressed block offset &= ~( 0xffffffffffffffffFULL  
<< compressed size bit shift )
```

Note that multiple compressed cluster data blocks seem to be always stored together in cluster block sizes. The compressed cluster data blocks are sector (512 bytes) aligned.

The compressed data uses a zlib inflate window bits value of -12

Compressed chunk data block – version 2

Version 2

```
compressed size bit shift = 62 - ( number of cluster block bits - 8 )
```

[MCLOUGHLIN08] describes the compressed block size calculation as:

```
compressed block size = ( ( ( cluster block offset & 0x3ffffffffffffffffFULL )  
>> compressed size bit shift ) + 1 ) * 512.
```

Since the compressed block size is stored in 512 byte sectors this value does not contain the exact byte size of the compressed cluster block data. It sometimes lacks the size of the last partially filled sector and one sector should be added if possible within the bounds of the cluster blocks size and the file size.

```
cluster block offset &= ~( 0xffffffffffffffffFULL  
<< compressed size bit shift )
```

Note that multiple compressed cluster data blocks seem to be always stored together in cluster block sizes. Although the file size does not seem to be strictly a multitude of the the cluster block size.

The compressed data uses a zlib inflate window bits value of -12

5. Snapshots

As of version 1 QCOW can use the backing filename in the file header to point to a parent image that contains the snapshot image where the current image only contains the modifications. Version 2 adds support to store snapshot inside the image.

5.1. Snapshot header - version 2

An in-image snapshot is created by adding a snapshot header, copying the L1 table and incrementing the reference counts of all L2 tables and data clusters referenced by the L1 table.

The snapshot header is variable of size and consists of:

offset	size	value	description
0	8		Level 1 table offset
8	4		Level 1 size

offset	size	value	description
12	2		Identifier string size
14	2		Name size
16	4		Date in seconds
20	4		Date in nano seconds
24	8		VM clock in nano seconds
32	4		VM state size
36	4		Extra data size
40	...		Extra data
...	...		Identifier string size
...	...		Name

TODO

6. Notes

6.1. Reference count table

reference count cluster block offset = cluster data block offset /
reference count table offset = cluster data block /

In order to obtain the reference count of a given cluster, you split the cluster offset into a refcount table offset and refcount block offset.

Since a refcount block is a single cluster of 2 byte entries, the lower cluster_size - 1 bits is used as the block offset and the rest of the bits are used as the table offset.

One optimization is that if any cluster pointed to by an L1 or L2 table entry has a refcount exactly equal to one, the most significant bit of the L1/L2 entry is set as a "copied" flag. This indicates that no snapshots are using this cluster and it can be immediately written to without having to make a copy for any snapshots referencing it.

Appendix A. References

[MCLOUGHLIN06]

Title: The QCOW Image Format
Author(s): Mark McLoughlin
Date: June 21, 2006
URL: <http://people.gnome.org/~markmc/qcow-image-format-version-1.html>

[MCLOUGHLIN08]

Title: The QCOW2 Image Format
Author(s): Mark McLoughlin
Date: September 11, 2008
URL: <http://people.gnome.org/~markmc/qcow-image-format.html>

[QEMU]

Title: QEMU Open Source Process Emulator
URL: http://wiki.qemu.org/Main_Page

Appendix B. GNU Free Documentation License

Version 1.3, 3 November 2008

Copyright © 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc.
<<http://fsf.org/>>

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 noncommercially. 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 in order 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.

The "publisher" means any person or entity that distributes copies of the Document to the public.

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 under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

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. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

11. RELICENSING

"Massive Multiauthor Collaboration Site" (or "MMC Site") means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A "Massive Multiauthor Collaboration" (or "MMC") contained in the site means any set of copyrightable works thus published on the MMC site.

"CC-BY-SA" means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

"Incorporate" means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is "eligible for relicensing" if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.