Windows Search forensics

Analyzing the Windows (Desktop) Search Extensible Storage Engine database

By Joachim Metz <jbmetz@users.sourceforge.net>

Summary

While some may curse Windows Vista for all its changes, for us forensic investigators it also introduced new interesting 'features'. For one is the integration of Windows (Desktop) Search into the operating system. Most corporations have been reluctant to adopt Vista, however more and more Windows XP systems are being replaced by Windows 7 equivalents. Windows 7 also contains Windows Search and enables by default is. It actually can be challenging to disable it. So one can conclude that Windows Search is becoming a relevant source of information in forensic analysis of Windows systems.

What is not widely known is that Windows Search uses the Extensible Storage Engine (ESE) to store its data. This is the same engine that Microsoft Exchange uses. Because ESE uses a propriety database format, little information about it is available in the public domain. As a consequence, it is unclear how well different forensic tools support the ESE database format.

After several years after the introduction of Windows Vista and Windows Search, currently only a handful of forensic analysis tools seem to provide support the Windows Search database. Although a Windows Search database can be a valuable source of evidence. This document provides an overview of the ESE database format and the Windows Search database and what it might contribute in your investigations.

Background

Although the Extensible Storage Engine (ESE) is a generic database engine, forensic analysis of ESE databases seem to be centered around Exchange. Little information about forensic investigation of ESE databases in general, seem to have been published in the public domain. As far as I can tell, Mark Woan author of EseDbViewer, was one of the first who published information about forensic analysis of ESE databases in general. This was in 2008.

Early 2009, I was getting search results in Windows.edb files (Windows Search databases) on Windows XP system in some investigations. Neither EnCase or FTK seem to offer any support for this file, although they claim to have EDB support. Not many other tooling seemed to be available to analyze the Windows Search ESE database. However when investigation Windows Vista system the Widows.edb file no longer contained any relevant results.

Besides trying to verify my assumptions on the Exchange related parts in the Microsoft Exchange OST files, this triggered me to start working on the ESE database format. I therefore started the libesedb project in September 2009. Findings from the libesedb projects and some of Mark Woan's EseDbViewer have been integrated in this document.

Document information

Author(s): Joachim Metz <jbmetz@users.sourceforge.net>

Abstract: This document contains information about the forensic analysis of the Windows Search Extensible Storage Engine database.

Classification: Public

Keywords: Windows Search, Microsoft Search, Extensible Storage Engine, ESE, EDB

License

Version

Version	Author	Date	Comments
1.0	Joachim Metz	June 2010	Initial version.

Table of Contents

1. Overview of the ESE database format....1
    1.1. Database header....1
    1.2. Page based storage....2
    1.3. Database tables and indexes....3
2. Analysis of a Windows Search database....4
    2.1. Data obfuscation....5
    2.2. Data compression....6
    2.3. Investigative artifacts and usefulness....6
    2.4. The Vista welcome mail....7
3. Conclusion....10
Appendix A. References....12
Appendix B. GNU Free Documentation License....13

1. Overview of the ESE database format

The Extensible Storage Engine (ESE) database format is mainly known for its use in the Microsoft Exchange, i.e. for the priv1.edb file. What is less widely known that a lot of Microsoft products use this file format, some of which are Active Directory (ntds.dit), Windows (Desktop) Search (Windows.edb) and Windows Mail (WindowsMail.MSMessageStore).

ESE is also known as Jet Blue in contrast to Jet Red that refers to the Microsoft Access database format. Microsoft has kept the specification of ESE database format closed, although the Jet Blue API has been partially documented on MSDN. The information in this document was obtained by the information available on the Internet and reverse engineering of the file format. The information obtained is maintained in a working documented titled: the Extensible Storage Engine (ESE) database (DB) format specification [ESEDB09].

There are three main variants of the ESE, one for Exchange 5.5 (ESE97), one for Exchange 2000 and later (ESE98) and one for Windows NT and later (ESENT). Active Directory and Windows Search use the ESENT version.

Basically an ESE database consists of the following elements:

- database header and a backup
- pages containing:
  - space tree data
  - database table data
  - database index data
  - long value data

The following paragraphs provide an overview of some of these elements.

1.1. Database header

The ESE database starts with a database header. The effective size of the database header is at least 667 bytes of size, e.g. the first 16 bytes.

```
00000000: 5c ca 88 0b ef cd ab 89  20 06 00 00 00 00 00 00   \.......  .......
```

Bytes 4 to 8 of the database header contain the unique signature '\xef\xcd\xab\x89' of the ESEDB format. Other significant values in the header are the file type, format version and revision and page size. The database header is actually stored in a block the size of a page; which is directly followed by another block containing a backup of the database header. This is one of the data redundancy measures provided in the ESE database format.

Different versions of Windows NT use different revisions of ESE, e.g. Windows XP uses version 0x620 revision 9, Windows Vista uses version 0x620 revision 12 and Windows 7 uses version 0x620 revision 17. Different revisions can have different methods of storing data, e.g. the Windows 7 version of ESE allows for 'native' compression of data; in previous versions applications using ESE needed to do compression themselves, like the (RTF) LZFu compression used by Exchange.

When no measures are taken to detect and handle compressed data, linear search and index-based search techniques will fail. So these techniques do not suffice for finding all the strings in ESE databases.

The ESE database format is also used for streaming file, e.g. priv1.stm used by Exchange, however until now little is know about the specifics of these streaming files. ESE uses transaction logs, which in theory could be used to analyze different versions of the data and mutations. However version analysis currently is in a state of infancy.

ESE comes with the eseutil (or its equivalent esentutl). Eseutil can be used to print the database header of an ESE database. The following example prints the database header of a Windows Vista Search (Windows.edb).

```
eseutil.exe /mh Windows.edb
```

```
Initiating FILE DUMP mode...
         Database: Windows.edb

        File Type: Database
   Format ulMagic: 0x89abcdef
   Engine ulMagic: 0x89abcdef
 Format ulVersion: 0x620,12
 Engine ulVersion: 0x620,12
Created ulVersion: 0x620,12
```

Sometimes you can come across a 'dirty' database. This is a database that was not neatly closed. The following information in the header information will indicate if an ESE database is considered 'dirty'.

```
            State: Dirty Shutdown
```

A 'dirty' database can be repaired using the repair option in eseutil.

```
eseutil.exe /r Windows.edb
```

Repairing an ESE database will alter the database file, but might be necessary for tools that cannot open 'dirty' databases. Sometimes it is also necessary repair before eseutil can perform certain operations on 'dirty' databases. Note that a successful repair is not guaranteed. Libesedb [ESEDB09] will try to open the database in its 'dirty' state.

1.2. Page based storage

At the lowest level an ESE database stores its data in pages. The size of the pages is stored in the database header and is applied to the entire database. A single page consists of a header, values and an index. A page does not need to be entirely filled, therefore a page has 'page unallocated space' which can contain remnant data. This remnant data can be of interest for forensic analysis.

A feature of impact on this remnant data is 'ESE (page) zeroing' which overwrites unused pages with various byte values. The 'zeroing' can be performed manually, by eseutil, or automatically, during online backup. For Exchange online backup is controlled by the following Registry key.

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSExchangeIS\ParametersSy
stem\Zero Database During Backup
```

Currently the actual impact of ESE (page) zeroing for forensic investigations is unknown.

As of Windows Vista Seach, a page can contain an error correcting code (ECC). The Microsoft

documentation states these ECC can only recover single-bit errors. The actual ECC method is not documented. In Windows 7 three additional ECCs were added, which probably allows for multi-bit recovery. This is another data redundancy measure provided in the ESE database format. Note that libesedb currently does not corrects errors using ECCs.

A page can contain multiple page values. Eseutil can be used to print the page values in page. The following example prints the values in page 13 of a Windows Vista Search ESE database (Windows.edb).

```
eseutil.exe /m /p13 Windows.edb
```

```
Initiating FILE DUMP mode...
      Database: Windows.edb
          Page: 13

          expected checksum   = 0x5c54a3ab36656192
                    new checksum format
                              expected ECC checksum = 0x5c54a3ab
                              expected XOR checksum = 0x36656192

                         checksum <0x00FE0000,  8>:  6653122505280414098
(0x5c54a3ab36656192)
                     dbtimeDirtied <0x00FE0008,  8>:  4646
(0x0000000000001226)
                         pgnoPrev <0x00FE0010,  4>:  0 (0x00000000)
                         pgnoNext <0x00FE0014,  4>:  14 (0x0000000e)
                         objidFDP <0x00FE0018,  4>:  2 (0x00000002)
                           cbFree <0x00FE001C,  2>:  3636 (0x0e34)
                cbUncommittedFree <0x00FE001E,  2>:  0 (0x0000)
                        ibMicFree <0x00FE0020,  2>:  5151 (0x141f)
                     itagMicFree <0x00FE0022,  2>:  74 (0x004a)
                          fFlags <0x00FE0024,  4>:  10242 (0x00002802)
                    Leaf page
                    Primary page
                    New record format
                    New checksum format


TAG   0    cb:0x000d    ib:0x0000    offset:0x0028-0x0034    flags:0x0000
TAG   1    cb:0x0037    ib:0x000d    offset:0x0035-0x006b    flags:0x0004 (c)
TAG   2    cb:0x0033    ib:0x0044    offset:0x006c-0x009e    flags:0x0006 (cd)
...
TAG  73    cb:0x0057    ib:0x1025    offset:0x104d-0x10a3    flags:0x0005 (cv)
```

First the information about the page header is provided followed by locations of the page values. Each page value is defined a tag (or index entry) and controlled by three flags, which are identified by the characters c, d and v. The actual meaning of the flags is undocumented but the d-flag seems to be used for deleted or defunct values. These deleted values are not overwritten and therefore can be interesting from an investigative point-of-view.

Eseutil does not provide means to access the data in the page values, except for some database metadata tables, like the catalog and the space tree

1.3. Database tables and indexes

The definition of the database tables and indexes are stored in a table referred to as the catalog.

The name of this table is 'MSysObjects'. Each ESE database contains a catalog and its backup named "MSysObjectsShadow'.

The data of tables and indexes are stored in a hierarchy of pages (or page-tree). These page-trees are traversed by means of (page) keys.

Eseutil can be used to print table information, e.g. the table information of the 'MSysObjects' table of a Windows Vista Search ESE database (Windows.edb).

```
eseutil.exe /mm /tMSysObjects Windows.edb
```

```
Initiating FILE DUMP mode...
         Database: Windows.edb

****************************** META-DATA DUMP ******************************
Name                                                    Type    ObjidFDP    PgnoFDP
================================================================================
Windows.edb                                              Db            1          1

MSysObjects                                              Tbl           2          4
  Name                                                   Idx           4          7
  RootObjects                                            Idx           5         10
********************************************************************************
```

From the output we can learn that the 'MSysObjects' table has two corresponding indexes: 'Name' and 'RootObjects'. 'ObjidFDP' refers to an unique 'object' identifier for each table or index. 'PgnoFDP' contains the page number of the Father Data Page (FDP), which basically is the root page of the page-tree.

Eseutil can be used to print all the tables and indexes of the database.

```
eseutil.exe /mm Windows.edb
```

The libesedb project comes with a tool called esedbinfo which does a similar print all of the tables and indexes in the database.

For some tables eseutil will print a line containing 'Long Values'.

```
SystemIndex_0A                                           Tbl          21       1125
  <Long Values>                                          LV          261        743
```

Long values are used by ESE to store 'large' amount of data in a separate page values; in effect also a separate page-tree. According to [MSDN]:

> *ESE stores the long value separated if it is larger than 1024 bytes or if the record would not fit on a single database page if stored in record.*

2. Analysis of a Windows Search database

Windows Search stores its data in a file named:

```
%Profiles%/All Users/Application Data/Microsoft/Search/Data/Applications/Windo
ws/Windows.edb
```

Note that '%Profile%' is dependent on the Windows version. To access the Windows.edb file the the Windows Search service needs to be deactivated and the necessary access rights are required.

If the database is in a 'dirty' state it might be necessary to copy the transaction logs as well. According to Mark Woan, author of EseDbViewer, copying the entire Windows Search application directory often does the trick.

Access to the ESE database format is only a small step closer to the information in a Windows Search database. As far as I know, forensic tools like EnCase or Forensic Toolkit do not support the Windows Search database; although they support some types of ESE databases. Additional specialized investigative tools like Windows Search Index Examiner or EseDbViewer are necessary; at least EseDbViewer directly uses the ESE. You could also consider to write a tool for a quick-and-dirty export of the values in the tables using ESE yourself.

From a forensic point of view using ESE is not the preferred method, because the engine alters the data; at minimum ESE sets the database state to 'dirty'. However ignoring possible evidence is not an option either. Another issue is that ESE will not open 'dirty' databases.

The approach of exporting data directly from a Windows XP Search database works fairly well. However when it comes to Windows Vista you're out of luck. Most of the columns have changed from the text to a binary format. Also the binary data in these columns is no longer readable; they have been compressed and obfuscated. Windows 7 Search uses native ESE compression and has largely switched back to text columns again.

One of the more interesting columns 'System_Search_AutoSummary', which contains part of the content of an indexed item, is compressed and obfuscated in the XP, Vista and 7 versions of Windows Search.

2.1. Data obfuscation

According to [TECHNET]:

> *Index files are lightly obfuscated.*
>
> *If the obfuscation is removed, meaningful data from documents can be extracted. The data structures of the index files do not lend themselves to easy reconstruction of a complete document. However, someone with enough tenacity and time could reconstruct the text for the majority of a document.*

Actually the obfuscation method is fairly straight forward. The obfuscation method uses a XOR with a bitmask based on the location of the byte in the data and an initial 32-bit bitmask.

The initial bitmask is created by a 32-bit XOR of the values in the Windows NT security identifier (SID):

```
S-1-5-12
```

The SID is stored as the following byte values:

```
01 01 00 00 00 00 00 05 12 00 00 00
```

This results in a 32-bit bitmask of:

```
0x05000113
```

The data is obfuscated using a method similar to the one below.

```
bitmask32  = 0x05000113;

bitmask32 ^= (uint32_t) encoded_data_size;

for( encoded_data_iterator = 0;
     encoded_data_iterator < encoded_data_size;
     encoded_data_iterator++ )
{
          switch( encoded_data_iterator & 0x03 )
          {
                    case 3:
                    bitmask = (uint8_t) ( ( bitmask32 >> 24 ) & 0xff );
                    break;

                    case 2:
                    bitmask = (uint8_t) ( ( bitmask32 >> 16 ) & 0xff );
                    break;

                    case 1:
                    bitmask = (uint8_t) ( ( bitmask32 >> 8 ) & 0xff );
                    break;

                    default:
                    bitmask = (uint8_t) ( bitmask32 & 0xff );
                    break;
          }
          bitmask ^= encoded_data_iterator;

          data[ data_iterator++ ] = encoded_data[ encoded_data_iterator ]
                                    ^ bitmask;
}
```

2.2. Data compression

Windows Search compresses the data before obfuscating it. For this it uses multiple compression methods. All these compression methods and obfuscation correction are incorporated in the function 'MSSUncompressText' stored in a Windows Search specific DLL. The name of the DLL differs per version of Windows Search. A quick-and-dirty approach could be to call the function directly to decompress the binary data.

Some of the obfuscation correction and decompression techniques have been integrated into esedbexport which is included in libesedb project [ESEDB09]. For a Windows Search database esedbexport tries to convert the compressed values it knows about. Note that the libesedb project is still in alpha status and you might want to validate findings, if possible, with other tools.

2.3. Investigative artifacts and usefulness

So what makes the Windows Search database so interesting for forensic analysis? For starters the Windows Search database contains a table named 'SystemIndex_0A' which contains vast amount

of values about various of artifacts found on a Windows system, e.g. files and directories, emails, appointments, attachments, images, audio and video, Microsoft Internet Explorer (MSIE) history, etc.

Better yet, on Windows Vista and 7, Windows Search is activated by default running as a system service, silently collecting this data on the background. Most users will be totally unaware that Windows Search is actually indexing potential evidence; talk about a system ready for investigation.

A Windows Search database can contain metadata and partial content data of deleted files. For now it is unknown how long Windows Search retains its data. From personal experience I can say that a Windows Search database on my test system still contained metadata about a file I thought I had thoroughly erased from that system a half year before.

Windows Search also can index items from other sources like an Exchange sever; yet another location to find (deleted) emails.

2.4. The Vista welcome mail

To give an idea of the values in a Windows Search database consider the Windows Vista Mail welcome email message.

Column	**Value**
DocID	60
__SDID	75
System_Search_GatherTime	Jun 14, 2008 06:00:53 (UTC)
System_Search_Rank	707406378
System_Size	24374
System_FileAttributes	32
System_DateModified	Jun 14, 2008 06:00:47 (UTC)
System_DateCreated	Jun 14, 2008 06:00:17 (UTC)
System_DateAccessed	Jun 14, 2008 06:00:17 (UTC)
System_ItemTypeText	Windows Mail E-Mail
System_ItemType	.eml
System_ItemFolderPathDisplayNarrow	Inbox
System_FlagStatusText	Not flagged
System_ItemName	270729C9-00000001.eml
System_ItemParticipants	"Microsoft Windows Mail Team" <msoe@microsoft.com>
System_ItemFolderPathDisplay	Inbox
System_FileOwner	host\user
System_FileFRN	1970324836988523
System_ItemPathDisplay	Inbox\Welcome to Windows Mail

Column	Value
System_FileName	270729C9-00000001.eml
System_MIMEType	message/rfc822
System_IsRead	FALSE
System_Search_AutoSummary	*<See text box below>*

```
(Please do not reply to this message)

WELCOME TO WINDOWS MAIL

Windows Mail is the successor to Outlook Express

Windows Mail builds on the foundation of Outlook Express, adding a variety of
new features designed to make your e-mail experience more productive and fun,
while helping to reduce risks and annoyances such as phishing and junk e-mail.


GETTING STARTED
If you're upgrading from Outlook Express, Windows Mail can import your
existing account information and e-mail addresses. The first time you start
Windows Mail, you will be prompted to set up an e-mail account. If you skip
this step and want to set up a new account later, click the Tools menu, click
Accounts, and then click Add.

In addition to sending and receiving e-mail, you can use Windows Mail to read
newsgroups, which are Internet discussion forums where groups of people gather
to talk about common interests. To participate in a newsgroup (you can send a
message or just read what other people are talking about), click Microsoft
Communities in the folder pane. You can choose from a variety of newsgroups
devoted to Windows and other Microsoft products.

To get help using Windows Mail, click the Help menu, and then click View Help.
You can also get help from other Windows Mail users in the
microsoft.public.windows.vista.mail newsgroup.


NEW FEATURES

Improved e-mail searching
 * To quickly search your messages  in Windows Mail, you can type complete or
partial words into the search box. You'll instantly get a list of all of the
messages that contain those words. The list of results will show messages that
contain your search criteria in  both the headers and message text of your
mail messages.
 * For fast access to search, press CTRL+E to move the cursor into the search
box. Press ESC to clear the search box.
 * You can also search your e-mail inbox from Windows by using the search box.
Searching from Windows instead of Windows Mail will produce the same results:
matches are based on both the headers and message text of the mail in your
inbox.

Junk e-mail and phishing filters
 * Windows Mail now includes Microsoft SmartScreen technology to help keep
unwanted junk e-mail out of your Inbox. Suspected junk e-mail messages are
automatically moved to the Junk E-mail folder.
 * The anti-phishing features in Windows Mail help protect against phishing
```

```
messages, which attempt to trick you into revealing personal or financial
information. When Windows Mail detects a possible phishing message, it allows
you to view the message, but it blocks any links or dangerous content that
might be in the message. You can choose to delete a message, or to allow a
message that you know is safe.

Communities
 * Windows Mail Communities let you rate the usefulness of newsgroup messages
by clicking the Rate this Post button. This makes it easier and faster to find
helpful, trusted information in busy newsgroups.
 * The Communities rating feature uses Windows Live ID to help ensure that the
people who post messages in newsgroups are who they claim to be. (You can
still utilize the Communities feature without using Windows Live ID.)


ABOUT NEWSGROUPS
Windows Mail is about more than just e-mail. You can use Windows Mail to
access Microsoft's Help newsgroups at msnews.microsoft.com by clicking
Microsoft Communities in the folder pane. These newsgroups allow you to ask
questions and read answers from other people who are also using Microsoft
products.

What you should know before you get started

1. Find the appropriate group. You'll find  newsgroups covering most Microsoft
products. Picking the appropriate newsgroup is the best way to receive the
information you want. Select folders related to the product that you have
questions about. For example, the group "microsoft.public.powerpoint" would be
the plac
```

Column	**Value**
System_IsShared	FALSE
System_Importance	3
System_ParsingName	270729C9-00000001.eml
System_FlagStatus	0
System_Photo_ExposureTime	0
System_Message_FromAddress	<msoe@microsoft.com>
System_SFGAOFlags	1077936453
System_Message_FromName	Microsoft Windows Mail Team
System_ItemDate	Jun 14, 2008 06:00:17 (UTC)
System_Message_Store	Windows Mail
System_Title	Welcome to Windows Mail
System_Photo_FNumber	0
System_Message_ToAddress	<msoe@microsoft.com>
System_Subject	Welcome to Windows Mail
System_Message_ToName	New Windows Mail User
System_Author	"Microsoft Windows Mail Team" <msoe@microsoft.com>

Column	Value
System_Photo_ShutterSpeed	0
System_Message_DateSent	Jun 14, 2008 06:00:17 (UTC)
System_Photo_Aperture	0
System_Image_HorizontalResolution	0
System_Identity	0
System_Image_VerticalResolution	0
System_ThumbnailCacheId	11060607374901453619
System_Message_HasAttachments	FALSE
System_ComputerName	host
System_Document_DateCreated	Jun 14, 2008 06:00:18 (UTC)
System_Kind	email; communication
System_Document_DateSaved	Jun 14, 2008 06:00:48 (UTC)
System_ItemNameDisplay	Welcome to Windows Mail
System_ItemPathDisplayNarrow	Welcome to Windows Mail (Inbox)
System_ItemAuthors	"Microsoft Windows Mail Team" <msoe@microsoft.com>
System_DateImported	Jun 14, 2008 06:00:18 (UTC)
System_Search_Store	file
System_Photo_FocalLength	0
System_FileExtension	.eml
System_IsAttachment	FALSE
System_ImportanceText	Normal
System_KindText	E-mail; Communication
System_ItemUrl	file:c:/Users/user/AppData/Local/Microsoft/Windows Mail/Local Folders/Inbox/270729C9-00000001.eml

As you can see metadata and part of the content of the Welcome email have been stored in the Windows Search database.

3. Conclusion

In short Windows Search can be a valuable source of investigative information and as of Vista it is available by default.

Windows Search uses the Extensible Storage Engine (ESE) database format to store its data. Although the ESE database format is complex and still evolving, the means to access ESE databases are readily available on a Windows system.

Windows Search uses both compression and obfuscation. Therefore investigative methods like

linear and index-based searches will fail unless a tool has support for the Windows Search database, which currently not many investigative tools seem to have. The compression and obfuscation can be easily taken care of by using Windows Search own decompression function.

The next time you're analyzing a Windows system have a look at the Windows Search database, perhaps it will help you in solving your case.

Appendix A. References

[ESEDB09]
Title: Extensible Storage Engine (ESE) database (DB)
Author(s): Joachim Metz
URL: https://libesedb.sourceforge.net/

[MSDN]
Title: Microsoft Developer Network
URL: http://msdn.microsoft.com/

[TECHNET]
Title: Windows Indexing Features
URL: http://technet.microsoft.com/en-us/library/dd744700%28WS.10%29.aspx#WS_IndexingOutlookandExchange

[WOAN08]
Title: EseDbViewer
Author(s): Mark Woan
URL: http://www.woanware.co.uk/esedbviewer

Appendix B. GNU Free Documentation License

Version 1.1, March 2000

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other written 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. Secondarily, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License 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 that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you".

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. (For example, 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.

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 "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, whose contents can be viewed and edited directly and 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 has been designed to thwart or discourage subsequent modification by readers is not Transparent. 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 designed for human modification. Opaque formats include PostScript, PDF, 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 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.

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 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 publicly-accessible computer-network location containing a complete Transparent copy of the Document, free of added material, which the general network-using public has access to download anonymously at no charge using public-standard network protocols. 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 less than five).
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", and 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. In any section entitled "Acknowledgements" or "Dedications", preserve the section's title, 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 as "Endorsements" or to conflict in title with any Invariant Section.

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.

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, does not as a whole count as a Modified Version of the Document, provided no compilation copyright is claimed for the compilation. Such a compilation is called an "aggregate", and this License does not apply to the other self-contained works thus compiled with the Document, on account of their being thus compiled, if they 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 quarter of the entire aggregate, the Document's Cover Texts may be placed on covers that surround only the Document within the aggregate. Otherwise they must appear on covers around 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 provided that you also include the original English version of this License. In case of a disagreement between the translation and the original English version of this License, the original English version will prevail.

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.