Expert Witness File Format Specification

Revised: April 7, 2002

Introduction

Developed by ASR Data, the Expert Witness file format (aka E01 format aka EnCase file format) is an industry standard format for storing "forensic" images.
The format allows a user to access arbitrary offsets in the uncompressed data without requiring decompression of the entire data stream.
The specification does **NOT** provide for quantifyable assurance of integrity, it is up to the implementation to provide meaningful authentication
for **any** data contained in an "evidence file".

Overview

The Expert Witness Compression format can store a single image in one or more segment files. Each file consists of a standard 13-byte header, followed by a series of *sections*. The sections are typically arranged back-to-back. A section cannot span two files.

- A *Chunk* is a 32k run of data (64 standard sectors).
- All offsets are relative to the beginning of the segment file, unless otherwise noted.

File Header

Each file begins with the following 13-byte header. (This is not to be confused with the *header section*, below.)

Signature Part (8 bytes)...

Bytes:	3	1	1	1	1	1
Data:	"EVF"	0x09	0x0d	0x0a	0xff	0x00

Fields Part (5 bytes)...

Bytes:	1	2	2
Data:	0x01	1 or higher	0x0000
Meaning:		Segment Number	

The *Section*

Every *section* begins with the same standard data, with the following layout.

Offset:	Bytes:	Data:	Meaning:
0 (0x0)	16	"volume", "header", etc.	Section type string
16 (0x10)	8		64-bit offset in current file to the next section
24 (0x18)	8		64-bit byte-size of the section
32 (0x20)	40	0x00...	Padding
72 (0x48)	4		CRC of all previous setion data

Section Types

Expert Witness Compression uses the following section types: `header`, `volume`, `table`, `next`, and `done`. Some of these section types have unique data that begins directly after the standard section structure above.

'header' section...

Offset:	Bytes:	Data:	Meaning:
76 (0x4c)	to end of section	zlib compress()'ed data	Comments structure (see below)

Comment structure is simply a text string in the following tab- and newline-delimited format. (The data in each cell is separated by a tab character, and each row is separated by a newline character.) The first three lines are standard and must not change. The characters in the third line serve as reminders for the content of the fields in the fourth line. (The fourth line is the only line that needs to be customized.)

1								
main								
c	n	a	e	t	m	u	p	r
Case Number	Evidence Number	Unique Description	Examiner Name	Notes	Acquired Date	System Date	pwhash	char

- *Case Number, Evidence Number, Unique Description, Examiner Name, and Notes are free-form (provided they don't contain tab or newline characters).*
- *Acquired Date and System Date are in the form of: "2002 3 4 10 19 59" (March 4, 2002 10:19:59).*
- *pwhash should simply be the character '0'.*
- *char should be the one of these three characters: 'b', 'f', or 'n'. This represents "best", "fastest", or "no compression". Expert Witness Compression uses 'f'.*

The header section should appear in the first segment file only.

'volume' section...

Offset:	Bytes:	Data:	Meaning:
76 (0x4c)	4	1	Reserved
80 (0x50)	4		Chunk Count
84 (0x54)	4	64	Sectors per Chunk
88 (0x58)	4	512	Bytes per Sector
92 (0x5c)	4		Sector Count
96 (0x60)	20	0x00...	Reserved
116 (0x74)	45	0x00...	Padding
161 (0xa1)	5		Reserved (Signature)
166 (0xa6)	4		CRC of all previous 'volume' data, starting with offset 76

The volume section should appear in the first segment file only.

'table' section...

Offset:	Bytes:	Data:	Meaning:
76 (0x4c)	4	1	Chunk Count (for this table)
80 (0x50)	16	0x00...	Padding
96 (0x60)	4		CRC of all previous 'table' data, starting with offset 76
100 (0x64)	as long as necessary		Offset array (see below)
from end of offset array	to end of section		zlib compress()'ed data Chunks

The offset array is a series of back-to-back 4-byte unsigned integer values. Each entry is an offset to the start of a compressed 'Chunk'. The high bit of each value must be set! There must be one entry per Chunk.

Each table section can hold 16375 entries. If more entries are needed, you must create multiple table sections per file.

'next' and 'done' sections...
Each file ends with a 'next' or 'done' section. If the file is the last segment in an Expert Witness compressed image, the section will be named 'done'. Otherwise, the section will be named 'next' to indicate that another segment file must be read. The "next section" pointer for these section types points to the beginning of the section itself, since it is the last section in a file. These section types have no unique data.
