

Implementing BitLocker Drive Encryption for Forensic Analysis*

Jesse D. Kornblum, ManTech International Corporation
jesse.kornblum@mantech.com

Abstract

This paper documents the BitLocker Drive Encryption system included with some versions of Microsoft's Windows Vista. In particular it describes the key management system, the algorithms and modes used, and the metadata format. Particular attention is given to methods forensic examiners can use to access protected data. There are some unanswered questions about how the cryptosystem operates, including an undocumented key management decision. This decision could allow, in a particular usage scenario, unauthorized access to a protected volume.

Keywords: BitLocker, Encryption, Key Management, Windows Vista, Elephant

1 Introduction

Some versions of Microsoft's Windows Vista include a Full Volume Encryption feature called BitLocker Drive Encryption. This feature enabled users to encrypt the system volume in the original version of Windows Vista and additional volumes as of Service Pack 1. The BitLocker system was intended to be used with a Trusted Platform Module (TPM) chip on the computer's motherboard and provide strong but unobtrusive protection for data at rest.

In its default mode, BitLocker stores a series of keys on each protected volume and in the TPM.

*This is the author's version of a work that was accepted for publication in *Digital Investigation*. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. A definitive version was subsequently published in *Digital Investigation* and is available at <http://dx.doi.org/10.1016/j.diin.2009.01.001>.

When the system is booted, the integrity of the operating system and hardware is verified. If the verification succeeds, the TPM releases an encryption key that allows the system to continue booting. The user does not have to provide any information to decrypt the volume. If a protected volume is removed from the system, however, it may be difficult if not impossible for an examiner to read the protected data. Decrypting the data without the keys stored in the TPM is infeasible.

Other modes of BitLocker operation allow the system to require a PIN to be entered, a specific removable storage device to be connected, or both, for a protected volume to be unlocked. A protected volume can also be unlocked using a recovery password, or a 48 digit hexadecimal number typed by the user.

BitLocker can be disabled without the protected volume being decrypted. When disabled no authentication or TPM verification is needed for the volume to be accessed, but it remains encrypted. Instead Windows creates a new encryption key and writes it to the protected volume. As explained below, this key can be used to decrypt the existing series of keys necessary to access the protected data. Like leaving a house key under the doormat, the volume is still protected, but by knowing where to look it's trivial to bypass the protection.

The Microsoft Corporation has provided detailed documentation on BitLocker Drive Encryption. The cryptosystem is detailed in [1] and some of the key management in [4, 5, 6, 7]. Kumar and Kumar's paper and source code added more detail on the key management system [2].

Unfortunately these documents do not provide all of the information necessary to create tools for the forensic analysis of BitLocker protected vol-

umes. Although a BitLocker protected volume can be mounted on another computer running Windows Vista, the examiner may prefer or be required to use another operating system.

This paper contains details necessary to access BitLocker protected volumes but not included in any previously published documentation. It describes how BitLocker operates, what cryptographic primitives it employs, how those primitives are implemented, and how the keys for those primitives are stored. Section 2 gives details on how the data at rest is protected. Section 3 describes the means by which different keys are manipulated to access the data at rest. Sections 4 and 5 describe the metadata BitLocker maintains. Finally, some unanswered questions about the cryptosystem and how it operates are presented in section 6.

2 Cryptosystem Overview

The BitLocker cryptosystem was developed by Niels Ferguson and mostly relies on previously published cryptographic primitives [1]. The data on a BitLocker protected volume is encrypted in one of four methods, all of which use the Advanced Encryption Standard (AES) in Cipher Block Chaining mode (CBC). The user can configure whether to use the 128 bit or 256 version of AES as well whether or not to diffuse the encrypted data. The default mode is to use 128-bit AES with the diffuser enabled.

When data is encrypted it is first XOR'ed against a sector key, optionally diffused, and then encrypted with AES-CBC. Decryption is the reverse: The data is first decrypted using AES-CBC, optionally diffused, and then XOR'ed against a sector key. Note that the diffuser contains two functions, A and B. When decrypting they must be run in the reverse order, B then A.

All of the key material used to encrypt and decrypt data comes from the 512-bit Full Volume Encryption Key (FVEK). When working with 128-bit AES, bits 0-127 of the FVEK are used in the AES-CBC key and bits 256-383 are used in the sector key. The remaining bits are not used, as shown in Figure 1. When working with 256-bit AES, bits 0-255 of the

FVEK are used in for the AES-CBC key and bits 256-511 are used in the sector key, as shown in Figure 2.

The diffuser mentioned above, Elephant, is a non-standard cryptographic algorithm invented by Ferguson and added to the cryptosystem to provide “additional security properties that are desirable in the disk encryption setting but which are not provided by AES-CBC cipher methods” [1]. Specifically the diffuser was added to prevent a manipulation attack. In such an attack a malfeasant could change a small amount of ciphertext in the hopes of changing a small amount of plaintext, such as a security setting. By altering just a small piece of data, the malfeasant could weaken the security of the system. By diffusing any change in the ciphertext throughout the plaintext, such targeted manipulations are much more difficult. Although Ferguson described the Elephant diffuser in [1], he did not provide a reference implementation. The description is sufficient to implement the system, however, and Kumar and Kumar were gracious enough to do so in [2].

3 Key Management

The BitLocker key management system uses a series of keys to protect the data at rest. This section describes these keys as they have been documented by Microsoft. Additional details developed from reverse engineering the system are then presented to demonstrate how the key management system was implemented.

The key used to protect the data at rest, the Full Volume Encryption Key, is stored on the protected volume. To prevent unauthorized access the FVEK is encrypted using another key. In particular the FVEK is encrypted using a 256-bit AES key working in Counter with CBC-MAC (AES-CCM) mode. (Although the CCM standard was originally defined using 128-bit keys [9], Microsoft has extended it to 256-bit keys. It should also be noted that Ferguson, the designer of the BitLocker cryptographic system, was also a co-author of CCM mode.) The key used to encrypt the FVEK, the Volume Master Key or VMK, is also stored on the protected volume. In

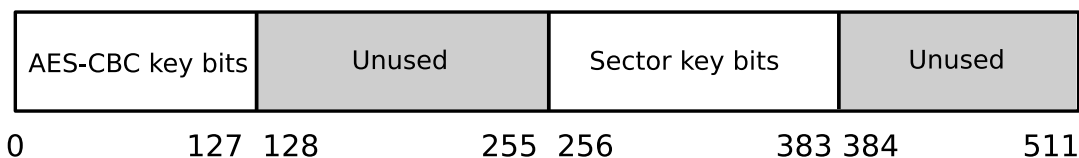


Figure 1: Layout of the FVEK in the 128-bit modes

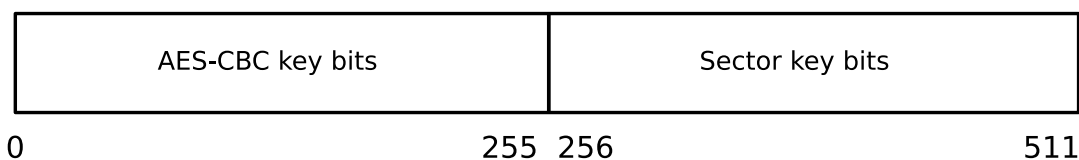


Figure 2: Layout of the FVEK in the 256-bit modes

fact several copies of the VMK are stored on the protected volume. Each copy of the VMK is encrypted using a different key. The different keys allow different access mechanisms to be used to access the stored data. Each access mechanism can be used to decrypt a copy of the VMK which in turn is used to decrypt the FVEK which in turn is used to decrypt the protected data.

The supported access mechanisms in Windows Vista, documented in [5], are a Trusted Platform Module (TPM) chip, TPM plus a PIN, TPM plus an external USB device (aka “Startup Key”), an external USB device, a recovery password, or an unprotected key saved to the protected volume. The last case is used to disable the BitLocker system without requiring the user to decrypt the protected data.

Using AES-CCM mode to encrypt these keys allows BitLocker to determine if a decryption operation has been successful. An AES transform can be applied to any data using any key and be considered “successful.” But because the keys being decrypted are indistinguishable from random data, the system cannot perform sanity checking on any computed plaintext. Thus the ciphertext stored with any data encrypted in AES-CCM mode includes a Message Authentication Code (MAC), or a kind of hash computed using the plaintext.

In accordance with the AES-CCM specification [9], a 16 byte MAC is computed using the plaintext. The

MAC of each key is stored along with the encrypted version of that key. When the system decrypts the ciphertext, it can compute a MAC for the decrypted value. If the computed MAC matches the stored MAC, then the system can assume, with an error probability of 2^{-128} , it has successfully decrypted the key.

The encryption and decryption operations of AES-CCM mode also depend on a nonce, or a number used once to generate an initialization vector. The nonce is also stored with the ciphertext.

At this point we begin describing how these cryptographic primitives are implemented in BitLocker. Specifically we are interested in how to conduct the decryption operations, where each of the above named values is stored, and any other information that can be gleaned from those values.

We’ll begin with the BitLocker key structure. This structure is used to hold encryption keys such as the FVEK. The structure, shown in Table 1, holds a size, two unknown values, the algorithm for which the key is intended, and the key itself. The data other than key add a total of twelve bytes to the structure. Thus the key structure for the 512-bit FVEK is 64 bytes of key plus 12 bytes of header for a total of 76 bytes.

When a key is encrypted the ciphertext is stored in a key protector structure as shown in Table 3. This structure contains a structure size, type, what BitLocker calls the Datum Type, version number, nonce,

MAC, and the ciphertext. The Datum Type in key protectors is normally five for what Microsoft labeled “AES-CCM”. The full list of datum type values, reverse engineered from fvevol.sys, is given in Table 4.

The nonce in each key protector is twelve bytes long. The first eight bytes are a FILETIME timestamp¹ of when the key protector was created. The last four bytes in the nonce are a counter value. Each key created by BitLocker gets a monotonically increasing counter value. The next counter value to be used is stored in the BitLocker metadata.

The nonces provide the examiner information on how many keys were generated that could be used to access the protected data. They also detail the order in which those keys were generated. That is, a key with a nonce of nine was generated after a key with a nonce of eight. The timestamps in the nonces indicate when, according to the system clock, each key was generated.

4 BitLocker Metadata

The data structures described above are stored on the protected volume in the BitLocker metadata. This section describes how an examiner can recognize a BitLocker protected volume, where to find the BitLocker metadata, and the format of the metadata header.

As noted in [8], a BitLocker protected volume is denoted by the string `-FVE-FS-` in lieu of the expected NTFS at offset three on the volume. If the BitLocker signature is found, the volume header should point to a block of BitLocker metadata. The offset for the first metadata block can be computed by multiplying the number of bytes per sector (found at offset 0xB in the volume header) by the number of sectors per cluster (offset 0xD) by the metadata Logical Cluster Number (LCN, offset 0x38). The value normally at offset 0x38, the LCN of the Master File Table (MFT) Mirror, is stored at offset 0x38 in the BitLocker metadata.

A BitLocker protected volume should contain three identical metadata blocks for redundancy [8]. The

volume’s metadata contains the offsets of all three metadata blocks, but the author is only aware of the first metadata block being listed at the start of the volume. If the start of the volume is damaged, the author is not sure how the operating system would find a metadata block. It is not apparent how the operating system resolves conflicts between the metadata blocks. That is, if two of the metadata blocks are identical but the third is different, it is unclear what the operating system does.

To find metadata blocks on a damaged volume, the examiner can search the volume for the metadata signature `-FVE-FS-`. Because each metadata block can only begin at offsets that are a multiple of the bytes per sector and sectors per cluster, the examiner could speed up the search by only searching for the string at these offsets. To be safe, the examiner should assume the smallest legal values and thus search for the BitLocker signature at multiples of 512 bytes.

Each BitLocker metadata block begins with a variable length header followed by a variable number of entries. A description of the first eight bytes of the BitLocker metadata header was given in [8], but had errors in the offsets. A more complete specification, based on the author’s reverse engineering and examination of protected volumes, is given in Table 5. The examiner should note that the validation data described in Section 5.6 is not included in the `Size` field as noted in [8].

There are several pieces of forensically valuable data in the header. First, the volume’s Global Unique Identifier (GUID) is stored at offset 0x50. This GUID should be included on any access device that unlocks this device such as USB sticks. Examiners can search for this GUID on USB devices to find possible BitLocker access devices. The date and time BitLocker was enabled is recorded at offset 0x68. Finally, the next counter value to be used for key encryption nonces is stored at offset 0x60. As mentioned earlier, this could be useful in determining how many access devices have been created for a volume.

¹Representing the number of 100-nanosecond intervals since 1 Jan 1601 (UTC)

Offset	Size	Field	Content
0x00	4	Size	
0x04	2	Unknown	
0x06	2	Unknown	
0x08	4	Algorithm	See Table 2
0x0C	Size - 0x0C	Key	

Table 1: BitLocker Key Structure

Value	Algorithm
0x1000	Stretch key
0x2000 - 0x2005	256-bit AES-CCM
0x8000	128-bit AES + Elephant
0x8001	256-bit AES + Elephant
0x8002	128-bit AES
0x8003	256-bit AES

Table 2: Encryption Algorithms

Offset	Size	Field	Content
0x00	2	Size	
0x02	2	Type	
0x04	2	DatumType	See Table 4
0x06	2	Version	Always one
0x08	12	Nonce	
0x14	16	MAC	
0x24	Size - 0x24	EncryptedData	Encrypted version of Table 1

Table 3: BitLocker Key Protector Structure

Datum Type	Description
0	Erased
1	Key
2	Unicode
3	Stretch Key
4	Use Key
5	AES-CCM
6	TPM Encoded
7	Validation
8	Volume Master Key
9	External Key
10	Update
11	Error

Table 4: Datum Types from fvevol.sys

Offset	Size	Field	Content
0x00	8	Signature	-FVE-FS- in ASCII
0x08	2	Size	Size of metadata, not including validation data
0x0A	2	Version	Should be one for both Vista RTM and Service Pack 1
0x0C	2	Unknown	
0x0E	2	Unknown	
0x10	16	Padding	All zeros
0x20	8	Metadata1	Offset of first metadata block
0x28	8	Metadata2	Offset of second metadata block
0x30	8	Metadata3	Offset of third metadata block
0x38	8	MFT Mirror	LCN of the MFT mirror
0x40	4	SizeMinusHeader	Size minus 0x40
0x44	4	Unknown	Always one
0x48	4	Unknown	Always 0x30
0x4C	4	SizeMinusHeader2	
0x50	16	Volume GUID	
0x60	4	NextCounter	Monotonically increasing counter for key protectors
0x64	4	Algorithm	Algorithm from Table 2 used to protect volume data
0x68	8	Timestamp	Time BitLocker was enabled in UTC
0x70	2	VolumeNameLength	
0x72	2	Unknown	
0x74	2	Unknown	
0x76	2	Unknown	
0x78	VolumeNameLength - 8	VolumeName	Volume's name in Unicode

Table 5: BitLocker Metadata Format

5 Metadata Entries

The BitLocker metadata header is followed by a series of metadata entries. These entries contain the encrypted FVEK and several encrypted copies of the VMK. Each copy of the VMK is encrypted with a different key. The keys used to encrypt the copies of the VMK, detailed in [4, 5], can be stored in the metadata, on an external device, entered by the user, released by the TPM, or a combination of these. If the system is configured to unlock the volume using a TPM chip, it encrypts that copy of the VMK using the RSA algorithm. For all other types of keys the VMK is encrypted using 256-bit AES-CCM. The remainder of this section describes the format of the individual metadata entries as determined through a combination of the literature and reverse engineering.

Each metadata entry consists of data concerning where the key in question is stored (e.g. in the TPM, on an external device) and at least two encrypted key protectors as described in Section 3. The first key protector structure contains a copy of the VMK encrypted with the key in question. That is, for the key k_i , there is a copy of the VMK encrypted with k_i , or $E(\text{VMK}, k_i)$. The second key protector contains a copy of the key in question encrypted with the VMK, or $E(k_i, \text{VMK})$. Although Kumar and Kumar claimed each metadata entry contained a copy of the key in question encrypted with itself, $E(k_i, k_i)$ [2], the author has found this not to be the case.

As noted earlier, each of these key protectors contains a nonce used in the encryption and decryption process. Those nonces increase monotonically for each key generation performed by BitLocker. The two nonces in the key protectors in each metadata entry differ by one and can be represented as j and $j+1$. The author has found that the nonce used when the system encrypts the key in question with the VMK, $E(k_i, \text{VMK})$, is the nonce j , and the nonce used when the system encrypts the VMK with the key in question, $E(\text{VMK}, k_i)$, is the nonce $j+1$. This suggests that when BitLocker creates a new metadata entry for a key, the system first encrypts the new key with the VMK and then encrypts the VMK with the new key. The timestamps are identical in both key protectors.

By storing a copy of each key k_i encrypted with the VMK and a copy of the VMK encrypted with each k_i , it is possible to use any valid key to recover all of the other keys. For example, let's assume that a protected volume has three metadata entries for three keys. Included in those entries are three copies of the VMK, each one encrypted with one of the individual keys. There are also encrypted versions of the three keys, each one encrypted with the VMK. Given any valid key, it is possible to decrypt the copy of the VMK encrypted with that key. At that point, given the VMK, it is possible to decrypt the other keys stored on the protected volume that were encrypted with the VMK. This issue and its implications for forensic analysis are discussed more in Section 6.

5.1 Test System

For this research the author used Windows Vista Ultimate Service Pack 1 to create a 10 MiB BitLocker protected volume. The author created an external key and a recovery password to unlock the volume. The author also configured the volume to “autounlock”, or disabled BitLocker. This created a clear key on the protected volume. These keys and metadata entries created for them are discussed in the following sections. All of the information in these subsections was determined through reverse engineering.

5.2 Full Volume Encryption Key

The encrypted FVEK was stored in the sample volume's metadata as shown in Figure 3. The entry begins with the entry's size, 0x70 bytes, in blue. The value three (3) that immediately follows indicates this is the FVEK. (A value of two (2) would indicate an encrypted copy of the VMK.) Next, the datum type of five indicates that the key is encrypted using AES-CCM. The nonce follows in red and then the MAC in green. As noted above, the nonce is a FILETIME timestamp of Tue Jul 1 16:05:11 UTC 2008 followed by a counter value. Finally, the encrypted data, an encrypted version of the key structure containing the FVEK, are shown in gray.

Offset	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
003661A0	70	00	03	00	05	00	01	00	70	67	60	3D	94	DB	C8	01
003661B0	08	00	00	00	D9	5C	82	49	D0	26	BB	89	28	20	33	E6
003661C0	A6	A1	3F	4A	E9	76	33	4B	4E	01	BB	A5	7F	3C	F4	9E
003661D0	C7	24	4E	DB	D0	3E	76	2A	4E	13	26	CB	7E	EF	5F	AF
003661E0	4F	77	91	87	F9	34	2D	89	8D	AC	71	37	AC	6A	FF	CC
003661F0	00	B8	D2	26	AB	15	39	DD	8A	75	D1	49	97	94	83	CD
00366200	B3	DB	01	F8	59	59	DA	D2	41	6D	FF	DF	8E	FD	46	E0

Figure 3: Encrypted FVEK in BitLocker metadata

5.3 External Keys

BitLocker can be configured to use a key placed on an external device such as a USB key. When using an external key, or to use Microsoft’s term, a Startup Key, the operating system stores a BitLocker External Key (BEK) on the device. For this research the author used the test system to generate an external key. The result was the BEK shown in Figure 4 and the metadata entry shown in Figure 5.

Both the BEK and the metadata entry contain the same Globally Unique Identifier (GUID). This allows the system to match the correct BEK to the correct metadata entry. The GUID, highlighted in orange, can be seen at offset 0x38 in the BEK and offset 0x366218 in the metadata entry. A key structure, starting at offset 0x70 in the BEK, contains the 256-bit key needed to decrypt the copy of the VMK stored in the metadata entry. The structure begins with the size of 0x2C bytes highlighted in blue. The type of data, a datum type from Table 4, is at offset 0x74. The key itself, highlighted in green, at offset 0x7C.

The metadata entry denotes that it is for an external key with the type value 0x20000000 at offset 0x366230, highlighted in light blue. The metadata entry contains two encrypted key protectors, starting at offsets 0x366260 and 0x3662B0 respectively. In each of these key protectors, the size is again highlighted in blue. The type of data, 0x5, again from Table 4, denotes AES-CCM encrypted data. Each key protector then has the identical time stamp highlighted in red (Tue Jul 1 16:13:39 2008 UTC), a MAC highlighted in green, and then the encrypted data highlighted in grey. As noted earlier, the first key pro-

tector contains the external key encrypted with the VMK and the second contains the VMK encrypted with the external key.

The external key, starting at offset 0x7C in the BEK, can be used to decrypt the data starting at offset 0x3662D0 in the metadata using the nonce starting at offset 0x3662B8. The result is the MAC starting at offset 0x3662C4 in the metadata and the key structure shown in Figure 6. If desired, the reader can verify that the VMK can be used to decrypt the other key protector (i.e. the data at offset 0x366280 using the nonce at 0x366268 to yield the MAC at 0x366274 and the external key at 0x7C in the BEK).

5.4 Recovery Password

When BitLocker is installed on a system volume the system encourages the user to create a recovery password. The user can also create a recovery password from the BitLocker Control Panel at any time. The recovery password can either be written directly to a text file or displayed on the screen. The “password” is actually a 48 digit number, eight groups of six digits, with three properties for checksumming described in [6, 7]. These properties are:

1. Each group of six digits must be divisible by eleven. This check can be used to identify groups mistyped by the user.
2. Each group of six digits must be less than 720,896, or $2^{16} * 11$. Each group contains 16 bits of key information. The eight groups, therefore, hold $16 * 8$ or 128 bits of key.

Offset	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
00000000	9C	00	00	00	01	00	00	00	30	00	00	00	9C	00	00	00
00000010	CF	34	72	9D	6D	00	6F	47	82	7B	17	9A	C0	32	D9	BE
00000020	01	00	00	00	00	00	00	00	F0	5C	3D	6C	95	DB	C8	01
00000030	6C	00	06	00	09	00	01	00	CD	2D	F7	17	42	38	A0	43
00000040	AF	23	73	05	9F	CA	2C	05	30	AD	20	6C	95	DB	C8	01
00000050	20	00	00	00	02	00	01	00	45	00	78	00	74	00	65	00
00000060	72	00	6E	00	61	00	6C	00	4B	00	65	00	79	00	00	00
00000070	2C	00	00	00	01	00	01	00	02	20	00	00	5A	84	D1	82
00000080	AA	05	B7	38	6C	4E	D7	B6	78	5A	BB	C9	1D	4D	AF	EF
00000090	EA	FA	66	31	F4	5D	44	0D	A5	DD	C4	B0				

Figure 4: Sample BitLocker External Key

Offset	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
00366210	F0	00	02	00	08	00	01	00	CD	2D	F7	17	42	38	A0	43
00366220	AF	23	73	05	9F	CA	2C	05	30	AD	20	6C	95	DB	C8	01
00366230	00	00	00	02	20	00	00	00	02	00	01	00	45	00	78	00
00366240	74	00	65	00	72	00	6E	00	61	00	6C	00	4B	00	65	00
00366250	79	00	00	00	5C	00	00	00	04	00	01	00	02	20	00	00
00366260	50	00	00	00	05	00	01	00	30	AD	20	6C	95	DB	C8	01
00366270	0D	00	00	00	2D	96	E7	84	41	D1	62	E6	AA	85	D3	7B
00366280	87	40	05	D0	67	00	68	71	9A	F6	D0	52	85	03	00	17
00366290	23	6D	7F	34	E9	24	C1	E7	80	C1	7A	CE	19	2E	AF	46
003662A0	24	FF	DB	3A	35	CB	72	4B	3F	EC	5F	6D	8E	C7	37	0A
003662B0	50	00	00	00	05	00	01	00	30	AD	20	6C	95	DB	C8	01
003662C0	0E	00	00	00	E9	5C	2B	DF	06	42	F4	19	9D	C0	D5	2A
003662D0	19	81	4D	D3	68	7A	7B	F4	E9	29	29	B7	9A	37	09	EA
003662E0	96	61	3F	5E	2D	CC	EC	7D	0A	09	F3	EA	AD	44	34	CC
003662F0	18	BE	25	EC	26	62	38	B5	26	3F	8A	4B	03	EB	54	27

Figure 5: Metadata Entry for External Key

Offset	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
00000000	2C	00	00	00	01	00	00	00	03	20	00	00	91	98	E3	96
00000010	2A	E0	7B	46	71	36	90	0B	0C	64	9A	E5	09	E8	8B	C1
00000020	62	56	DB	AC	AA	A4	A2	0E	6D	6C	06	07				

Figure 6: Decrypted VMK

- The sixth digit in each group is a checksum digit. If the digits in each group are represented as x_1, x_2, x_3, x_4, x_5 , and x_6 , then x_6 must equal $(x_1 - x_2 + x_3 - x_4 + x_5) \bmod 11$. Note that the Microsoft documentation for the check digit calculation is incorrect. In [7] Microsoft claimed that x_6 must equal $(-x_1 + x_2 - x_3 + x_4 - x_5) \bmod 11$. The error may be attributable to the fact that many programming languages do not compute the modulus operator correctly. In C, for example, `-5 % 11` yields `-5`, and not the correct result, `6`. A method to compute the check digit correctly in C is: `check_digit = (11 - (-x1 + x2 - x3 + x4 - x5)) % 11;`. The author believes that whomever wrote the documentation for Microsoft may have relied on the source code. If that source code resembles the version of the checksum above, it would account for the discrepancy in [7].

To create the metadata entry corresponding to a recovery key, BitLocker chooses a recovery key and a salt value. The recovery password is distilled down to a 128-bit value and then chain hashed with a salt value stored in the metadata to produce an “intermediate key”. This intermediate key is used to encrypt a copy of the VMK. The VMK is also used to encrypt a copy of the intermediate key. The metadata entry for a recovery key contains these two encrypted key protectors and the salt value.

The recovery password is distilled by dividing each group of six digits in the recovery password by eleven and then converting them to a series of little endian bytes. For example, the BitLocker recovery password groups 480370 and 051986 would be divided by eleven to become 43670 and 4726, or in hexadecimal, 0xAA96 and 0x1276. When converted to a series of little endian bytes, they become 0x96 0xAA 0x76 0x12.

The chain hashing procedure was described by Kumar and Kumar in [2] and confirmed by the author. The method uses an 88 byte structure, shown in Table 6, that holds a hash of the distilled recovery password, the salt from protected volume’s metadata, a counter, and an updated hash value. Initially the counter is set to zero. Over 2^{20} iterations, the entire structure is hashed using SHA-256 and the result

stored in the structure itself. The counter value is incremented and then the next iteration begins. An implementation for the chain hashing procedure in C is shown in Figure 7.

The author used the test system to produce a recovery password. The system produced the password 004301 051986 278476 162294 184228 193919 575828 424457 and created the metadata entry shown in Figure 8 starting at offset 0x3660AE. The entry contains an overall size value (blue), the GUID for the recovery password in orange, and the entry type, 0x80000000, highlighted in cyan. Immediately following the entry type is a Unicode text string, “Disk Password”. The author did not observe Unicode strings in the metadata of BitLocker protected volumes created with Windows Vista RTM.

Following these values, at offset 0x366100, is the salt to be used in chain hashing and two encrypted key protectors. Note that the first key protector contains 0x10 fewer bytes of encrypted data than the second. This is because the first key protector contains the distilled bytes from the recovery password encrypted with the VMK. That key, only 128-bits, takes up only 0x1C bytes. The second key protector contains a copy of the VMK encrypted with the intermediate key. The 256-bit VMK, inside of its key structure, uses 0x2C bytes.

Using the distillation method described above, the recovery password blocks from the test system (004301 051986...) were each divided by eleven, yielding 391 4726 25316 14754 16748 17629 52348 38587. When converted to a series of little endian bytes they become 0x87 0x01 0x76 0x12 0xE4 0x62 0xA2 0x39 0x6C 0x41 0xDD 0x44 0x7C 0xCC 0xBB 0x96. This 128-bit value is then chain hashed with the salt from Figure 8 at offset 0x366100 to produce the 256-bit intermediate key shown in Figure 9.

```

0x9F 0x44 0x31 0x30 0x8F 0xB1 0x1A 0xE3
0x4D 0xE4 0x19 0x8E 0x51 0x97 0x48 0x38
0xE1 0xD5 0xE5 0x00 0x0A 0xE3 0x8F 0xEF
0x30 0x89 0x82 0xFC 0xBA 0x70 0xF8 0xDE

```

Figure 9: Recovery Password Intermediate Key

```

typedef struct {
    /* 0x00 */ unsigned char  updated_hash[32];
    /* 0x20 */ unsigned char  password_hash[32];
    /* 0x40 */ unsigned char  salt[16];
    /* 0x50 */ uint64_t       hash_count;
} bitlocker_chain_hash_t;

// Chain hash the given 16 byte recovery key previously
// distilled from a 48 digit recovery password. Uses a 16
// byte salt value. Stores the 32 byte result in result.
int chain_hash(const unsigned char * recovery_key,
               const unsigned char * salt,
               unsigned char * result)
{
    sha256_context ctx;
    size_t size = sizeof(bitlocker_chain_hash_t);
    bitlocker_chain_hash_t * ch;

    ch = (bitlocker_chain_hash_t *)malloc(size);
    if (NULL == ch)
        return true;

    memset(ch,0,size);

    sha256_starts(&ctx);
    sha256_update(&ctx,recovery_key,16);
    sha256_finish(&ctx,ch->password_hash);

    memcpy(ch->salt, salt, 16);

    for (uint64_t loop = 0 ; loop < 0x100000 ; ++loop)
    {
        sha256_starts(&ctx);
        sha256_update(&ctx,ch,size);
        sha256_finish(&ctx,ch->updated_hash);

        ch->hash_count = ch->hash_count + 1;
    }

    memcpy(result,ch->updated_hash,32);

    free(ch);
    return false;
}

```

Figure 7: Chain Hashing Implementation in C

Offset	Size	Field	Content
0x00	32	UpdatedHash	SHA-256 hash result updated on each iteration
0x20	32	PasswordHash	SHA-256 hash of the distilled Recovery Password
0x40	16	Salt	Salt from the metadata
0x50	8	Counter	Iteration counter

Table 6: Chain Hashing Structure

Offset	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
003660A0	31	00	2F	00	32	00	30	00	30	00	30	00	00	00	F2	00
003660B0	02	00	08	00	01	00	F8	CB	B9	4C	A8	9E	0F	40	A2	CB
003660C0	39	14	DE	20	7A	44	20	9B	22	3D	94	DB	C8	01	00	00
003660D0	00	08	22	00	00	00	02	00	01	00	44	00	69	00	73	00
003660E0	6B	00	50	00	61	00	73	00	73	00	77	00	6F	00	72	00
003660F0	64	00	00	00	5C	00	00	00	03	00	01	00	00	10	00	00
00366100	3B	36	D9	30	72	A2	2E	03	F2	ED	FE	6F	CD	14	B4	58
00366110	40	00	00	00	05	00	01	00	60	6A	6B	3A	94	DE	C8	01
00366120	03	00	00	00	F9	C3	B5	61	B5	E6	B6	C1	19	3A	20	DC
00366130	CD	D9	32	27	E5	51	E0	0A	78	49	B6	C2	8B	10	75	14
00366140	5F	38	85	1D	16	65	1A	F0	15	63	7E	B7	DF	45	35	E1
00366150	50	00	00	00	05	00	01	00	60	6A	6B	3A	94	DE	C8	01
00366160	04	00	00	00	D4	07	C0	12	E2	EB	E7	1C	A5	A6	4D	EA
00366170	45	E1	61	9E	31	60	EC	C9	DD	86	2F	06	C2	68	84	AB
00366180	6A	0C	5C	D6	37	88	7B	19	2C	A1	71	06	E9	8D	6F	51
00366190	5F	4E	23	08	46	41	35	CC	19	42	8D	E8	87	05	D4	FE

Figure 8: VMK encrypted with Recovery Password

The computed 256-bit intermediate AES key can be used to decrypt the encrypted copy of the VMK at offset 0x366150 using AES-CCM. If successful, the decrypted data should contain the key structure containing the VMK shown in Figure 6. The reader can verify the VMK can also be used to decrypt the encrypted key protector at offset 0x366110 to yield the little endian bytes distilled from the recovery password (i.e. 0x87 0x01 0x76 0x12 0xE4...).

5.5 Clear Key

A user can disable BitLocker without decrypting the volume. This could be done to allow for the installation of new hardware or to disable any authentication checks. When BitLocker is disabled but not turned off, the operating system writes a 256 bit “clear key” to the volume’s metadata along with a copy of the VMK encrypted with that key. Thus the system can decrypt the VMK and FVEK without any other information.

The metadata entry for a clear key contains a GUID, a timestamp, the 256-bit key, and an encrypted version of the VMK. The VMK decryption procedure is the same as in the previous sections using the supplied key in AES-CCM mode.

Examiners should always check for the presence of a clear key when attempting to access a protected volume. If BitLocker has been disabled, accessing the drive should be easy. The examiner can use the clear key to decrypt the VMK and then use the VMK to decrypt the FVEK. Once the FVEK has been obtained, the examiner can decrypt the contents of the protected volume.

5.6 Validation Data

Following the metadata entries is a block of validation data. This block contains eight bytes of unknown data and an encrypted key structure. The encrypted data in the key structure is a SHA-256 hash of the metadata structure encrypted with the VMK, or $E(\text{SHA-256}(\text{metadata}), \text{VMK})$.

6 Unanswered Questions

There are several unanswered questions about BitLocker. First, it is unclear why Microsoft included a copy of each key encrypted with the VMK in the metadata entries. It could be intended as another MAC, just in case the MAC checks for decrypting both the VMK and FVEK fail. That is, both of those decryptions produce a value that is not the correct key, but the MAC matches the MAC for the correct key. Given that the probability of this happening is 2^{-256} , it is a rather unlikely occurrence. Granted, the consequences of booting a system using an incorrect key would be disastrous. At best the system should would crash, but at worst it could corrupt data on the protected volume. But the author does not know if that chance is worth the potential security risk of including the extra key.

Regardless of intent, having these extra keys could be helpful for forensic analysis. Importantly they could also be used by a malefactor to access data. Let’s assume that a protected volume is configured to use USB tokens for access. This computer is shared by several employees at an office. Each employee is given their own unique USB token. As noted above, when a user accesses the system with their token, the key on that token decrypts a copy of the VMK. The VMK can then be used to obtain the keys used on *all of the other tokens*. Any employee could obtain the keys used by other employees to access the system. If an employee’s access to the system was ever revoked, they could use one of those pilfered keys to create a new USB token and continue to access the system.

It can be argued that such an attack is unrealistic. If such a malefactor has access to the VMK then they also have access to the FVEK. Why should they bother creating a new USB token when they could use the VMK to capture the FVEK? They could then access the contents of the drive at will since the FVEK never changes. The attack described above, however, allows the user to access the protected volume without using any special tools. The malefactor, to outside appearances, would be using the system with “his” USB key. Other employees might be suspicious if they saw someone else accessing the system using a non-standard tool.

On the other hand, the extra keys in the metadata could be used by a forensic examiner to access a protected drive. If the examiner has one-time access to a system, she might be able to extract the other keys from the system and use them to generate her own access device for later use.

Although not an outright weakness in the cryptosystem, the author believes that either the encrypted copies of each key should be removed or that this behavior should be documented by Microsoft. As is, it represents an unnecessary security risk.

Next, the author is not aware of how the system resolves conflicts between metadata blocks. If one block is different from the others, which one is chosen to be “correct”? If the volume’s header becomes damaged, how does the system find any of the metadata blocks? Without the magic `-FVE-FS-` header, would it even know that BitLocker has been enabled? The unknown eight bytes in the validation data may contain valuable information. Resolving inconsistencies in damaged media is a primary concern for examiners. Even a small wiping operation could erase the disk header, leaving the system unable to find the BitLocker metadata information. Examiners should search for the `-FVE-FS-` signature at the start of each metadata block to find the keys on the damaged volume. In such a case a manual decryption might be the examiner’s only hope to access the protected volume.

Finally, the author has hypothesized that during normal operation, the system only keeps the FVEK in memory. This hypothesis is based on not seeing key schedules for VMK or the VMK itself in any memory images captured from the target system during normal operation [3]. Reverse engineering parts of the BitLocker system indicates that the system zeros out sensitive data as soon as they are no longer needed, consistent with good security practices. On the other hand, when the user requests a new key protector for the volume, the system must not only generate the key, but perform two operations with the VMK. Namely, it must encrypt the new key with VMK and encrypt the VMK with the new key. Where does the operating system get a copy of the VMK? Is there a copy of the VMK somewhere in memory? Are there other sensitive details in memory? Can forensic ex-

aminer find those data and exploit them?

7 Conclusion

The BitLocker Drive Encryption system provides protection for data at rest. The heart of the system, the 512-bit Full Volume Encryption key, is stored on the protected volume but encrypted using a series of keys. A forensic examiner can use a BitLocker access device to access the FVEK and thus the protected data. These access devices are themselves keys which can be used to decrypt the series of keys protecting the FVEK. Some pieces of the metadata surrounding these keys could be useful to a forensic examiner, including the order in which keys were generated, the number of keys generated, and the types of those keys. Additionally, some features of the key management system allows access to all of the access devices protecting a volume provided the user has a valid access device.

8 Acknowledgments

The author would like to thank all those who helped bring this paper together, including Elizabeth Lovegrove, Caroline Bauer, Robert J. Hansen, S—, WinHex, and caffeine.

References

- [1] Niels Ferguson. AES-CBC + Elephant diffuser A Disk Encryption Algorithm for Windows Vista. Technical report, Microsoft Corporation, September 2006.
- [2] Nitin Kumar and Vipin Kumar. Bitlocker and Windows Vista, May 2008. <http://www.nv1abs.in/node/9>.
- [3] ManTech International Corporation. *ManTech Memory DD*, 1.3 edition, August 2008. <http://mdd.sf.net/>.
- [4] Microsoft Corporation. BitLocker FIPS Security Policy, 2007. <http://csrc.nist.gov/groups/>

STM/cmvp/documents/140-1/140sp/140sp947.pdf.

- [5] Microsoft Corporation. Bitlocker drive encryption technical overview. Technical report, Microsoft Corporation, May 2008. <http://technet.microsoft.com/WindowsVista/en/library/ce4d5a2e-59a5-4742-89cc-ef9f5908b4731033.aspx?mfr=true>.
- [6] Microsoft Corporation. *ProtectKey-WithNumericalPassword Method of the Win32_EncryptableVolume Class*, February 2008. [http://msdn.microsoft.com/en-us/library/aa376467\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/aa376467(VS.85).aspx).
- [7] Microsoft System Integrity Team. Bitlocker recovery password details, August 2006. http://blogs.msdn.com/si_team/archive/2006/08/10/694692.aspx.
- [8] Microsoft System Integrity Team. Detecting bitlocker, October 2006. http://blogs.msdn.com/si_team/archive/2006/10/26/detecting-bitlocker.aspx.
- [9] D. Whiting, R. Housley, and N. Ferguson. Counter with CBC-MAC (CCM). RFC 3610 (Informational), September 2003.