Practical Cryptographic Data Integrity Protection with Full Disk Encryption

Milan Brož ^{1,2}, Mikuláš Patočka¹, and Vashek Matyáš²

¹ Red Hat Czech, Brno, Czech Republic
{mbroz,mpatočka}@redhat.com

² Masaryk University, Faculty of Informatics, Brno, Czech Republic
{xbroz,matyas}@fi.muni.cz

Abstract. Full Disk Encryption (FDE) has become a widely used security feature. Although FDE can provide confidentiality, it generally does not provide cryptographic data integrity protection. We introduce an algorithm-agnostic solution that provides both data integrity and confidentiality protection at the disk sector layer. Our open-source solution is intended for drives without any special hardware extensions and is based on per-sector metadata fields implemented in software. Our implementation has been included in the Linux kernel since the version 4.12.

1 Introduction

Storage encryption has found its way into all types of data processing systems – from mobiles up to big data warehouses. Full Disk Encryption (FDE) is a way to not only provide data confidentiality but also provide an easy data disposal procedure (by shredding the used cryptographic key). FDE and the underlying block device work with the disk sector as an atomic and independent encryption unit, which means that FDE can be transparently placed inside the disk sector processing chain.

A major shortcoming of current FDE implementations is the *absence of data integrity protection*. Confidentiality is guaranteed by symmetric encryption algorithms, but the nature of length-preserving encryption (a plaintext sector has the same size as the encrypted one) does not allow for any metadata that can store integrity protection information.

Cryptographic data integrity protection is useful not only for detecting random data corruption [27] (where a CRC-like solution may suffice) but also for providing a countermeasure to targeted data modification attacks [38]. Currently deployed FDE systems provide no means for proving that data were written by the actual user. An attacker can place arbitrary data on the storage media to later harm the user. FDE systems such dm-crypt [21, 4] or BitLocker [19] simply ignore the data integrity problem, and the only way to detect an integrity failure is the so-called *poor man's authentication* (user can recognize data from garbage produced by the decryption of a corrupted ciphertext).

The aim of our work is to *demonstrate that we can build practical cryptographic data integrity and confidentiality protection* on the disk sector layer with

acceptable performance and without the need for any special hardware. Our solution is an open-source extension of existing tools; we do not invent any new cryptographic or storage concepts. To achieve this goal, we implemented a per-sector metadata store over commercial off-the-shelf (COTS) devices. We provide an *open-source implementation as a part of the mainline Linux kernel*, where it is crucial to avoid proprietary and patented technology.

We focus on the *security of authenticated encryption and the algorithm-agnostic implementation*. Our main contributions are as follows:

- separation of storage and cryptographic parts that allow changing the underlying per-sector metadata store implementation without modifying the encryption layer,
- the concept and implementation of emulated per-sector metadata,
- algorithm-agnostic implementation of sector authenticated encryption in the Linux kernel and
- use of random initialization vector for FDE.

Storage security has often been perceived as an additional function that can easily be added later. The history of length-preserving FDE is a demonstration of this false idea. Although it is a simple application of cryptography concepts, some vendors deployed FDE not only with known vulnerabilities but also with incorrectly applied cryptography algorithms [7, 34].

The remainder of this paper is organized as follows. Section 2 discusses the threat model and introduces the data corruption problem. Section 3 describes proposed algorithms for encryption and integrity protection. Section 4 discusses how to store additional integrity data and extends the proposed metadata store concept to construct a reliable virtual device that handles both data and metadata. Sections 5 and 6 describe our practical implementation and performance evaluation. Section 7 concludes our paper and identifies future work.

2 Threat Model and Use Cases

Use cases for of FDE can be categorized to several situations, like *stolen mobile device or laptop*, *devices in repair*, resold *improperly wiped* devices, *virtual device* in a multi-tenant environment, or a *mobile device* storage. In all of these scenarios, confidential data can leak out of the control of the owner. A recent analysis of the content of used drives [29] shows that encryption is often still not used despite the importance of storage data encryption being recognized for a long time [32].

Our threat model adds to all these scenarios the detection of *data tampering* on leaked (not-wiped) devices and expects that an attacker has limited ability to record device changes with access snapshots of the device in time. The model recognizes the *Focused Opportunistic* attacker defined in Table 1. It does not protect data in situations when a device is in active (unlocked) state or an attacker can get the unlocking passphrase (and thus encryption key) directly. As mentioned in Section 1, our model expects COTS devices, it cannot rely on use

of any of tamper-proof cryptographic devices like Hardware Security Modules (HSMs).

2.1 Attackers

We define three simplified types of attackers, as summarized in Table 1. The most common type of attacker for FDE is a *random* attacker. We define a *focused opportunistic* attacker type for sophisticated attacks that focus on the situation where the stolen device is returned to the user. In the *targeted* attacker case, FDE will not provide sufficient protection without additional countermeasures.

Type	Description
Random	An attacker with user skills, can investigate disk content. Focuses on hardware, to steal and sell it. Valuable data are simply a bonus. A typical representative for this category is a random thief.
Focused Opportunistic	An attacker can run advanced forensic tools and can use documented attacks (such as a brute-force password search). The goal is to obtain data, analyze them, and use them to make a profit. The attack is not targeted to a specific user, but if the recovered data allow such an attack, the attacker can switch to this type of attack. A typical representative is a computer technician with access to devices for repair. In some cases, the attacker can access the device repeatedly, but only in limited opportunistic occasions.
Targeted	A top skilled attacker that uses the best possible attack vectors, typically to focus on a specific user with exactly defined goals in advance. The user system is generally under full attacker control.

Table 1: Attackers.

Type	Description
Pure FDE	Length-preserving encryption that provides confidentiality only.
Authenticated FDE	Encryption that provides both confidentiality and integrity protection, but limited by COTS devices (no hardware for authentication).
HW-trusted	The ideal solution with confidentiality and integrity protection. It stores some additional information to external trusted storage in such a way that the system can detect data replay.

Table 2: Discussed types of FDE protection.

2.2 FDE Protection Types

For the description of our model, we define three basic levels of protection, as summarized in Table 2. Here, FDE protection (of any type) means that the data confidentiality is enforced. A simple case of a device theft means that only hardware is lost. Data remain encrypted and not accessible to the attacker. This scenario is covered by the *Pure FDE* protection. The importance of *authenticated*

FDE comes into play when the stolen or seized device returns to the user (and this often occurs in reality; and example can be mandatory border checks).

This situation is generally enforced by a security policy and compliance of users. The reality is different – experiments show that people plug-in foreign devices, even if such devices are obtained under suspicious circumstances [37].

Authenticated encryption enforces that a user cannot read tampered data but will see an *authentication error*. It not only stops any attempts to use tampered data on higher layers, but also helps a user to realize that the device is no longer trustworthy. An overview of the features among FDE types is summarized in Table 3.

	None	Pure FDE	Auth. FDE	HW trusted
Confidentiality	×	✓	✓	✓
Integrity	×	×	✓	✓
COTS hardware	✓	✓	✓	×
Detect silent corruption	×	×	✓	✓
Detect data tampering	×	×	✓	✓
Detect data replay	×	×	×	✓
Whole sector change	×	(✓)	(✓)	(✓)

Table 3: Overview of FDE features.

2.3 Data Corruption and Forgery

So-called silent data corruption [9] is a common problem in persistent storage. This problem occurs when data are unintentionally and randomly corrupted while traversing through the storage stack. It is generally caused by flaky hardware (unstable memory bits and loose cables), triggered by an external influence (ionizing radiation) or by misplacement of data (correct data are written to an incorrect place). Data are then stored and used in a corrupted form.

A solution is to detect data corruption by checksums such as CRC32 [35]. This solution is not adequate if we want to detect an intentional unauthorized change (an attacker will simply fix a checksum).

An active attacker can not only cause undetected data corruption by simulating silent data corruption but can also attempt to forge the data by corrupting a specific disk area. A more sophisticated attacker can store additional data using steganographic techniques (conceal other data) to unused areas of a disk. We have to use a cryptographic integrity protection to detect such a situation.

2.4 Replay Attacks

In the strong view of cryptographic integrity protection, we should also detect data replacement using old content (revert to snapshots, also called *replay attack* [15]). Such a requirement cannot be fulfilled without an additional trusted metadata store independent of the storage itself. The content of the entire storage

can always be completely replaced with an older snapshot. Without additional and trusted information, users cannot recognize such a situation. An attacker can also revert only a partial part of the storage (in our case, selected sectors). From the cryptographic perspective, this situation cannot be completely prevented or detected (it would require breaking the independence of sectors).

In this text, we present algorithms that *do not protect* from the replay attack. This decision is based on the fact that our work is focused on utilizing standard disk drives without any additional hardware requirements.

3 Length-Preserving and Authenticated Encryption

Mode	Description
AES-CBC	AES [1] non-authenticated mode used in legacy FDE systems. [21]
AES-XTS	AES [1] non-authenticated mode used in recent FDE systems. [18]
AES-GCM	AES [1] in the Galois/Counter authenticated mode. [17] Due to only 96-bit nonce, it can be problematic in the FDE context. [12]
ChaCha20-Poly1305	Authenticated mode based on the ChaCha20 [11, 30] cipher and the Poly1305 authenticator.

Table 4: Examples of encryption algorithms.

Length-preserving encryption algorithms are used in current FDE solutions. These algorithms transform original data (plaintext) to its encrypted form (ciphertext), and confidentiality is based on the use of secret keys. Storage data are (for performance reasons) always encrypted with symmetric block ciphers.

The block sizes of these symmetric ciphers are typically 16 bytes. The device sector size is considerably larger (at least 512 bytes); thus, to apply encryption per sector,

we have to utilize a *block encryption mode* inside a sector. Current FDE systems use CBC [19] or XTS [18] modes. The CBC mode has many known problems [21] and specific requirements (see the note in Section 3.2).

The XTS mode, due to the internal parallel processing of blocks, can leak more information about a change in the plaintext than other modes. If only a single byte changes in a sector, then we can localize the change with the internal cipher block granularity. An ideal FDE system should produce a pseudo-random change of the entire sector data. Also these modes produce the same plaintext data encrypted to the same sector always produce the same ciphertext. In cryptography, this means that such a system does not fully provide indistinguishability under chosen plaintext attack (IND-CPA) [22, 26].

3.1 Authenticated Encryption

We have two options for integrity protection combined with device encryption: either to use *Authenticated Encryption with Additional Data* (AEAD) [10, 33] or to combine length-preserving encryption with an additional cryptographic integrity

operation. The major difference is that for the combined mode, we can ignore integrity tags and decrypt the data without such tags. In the AEAD mode, the authentication is an integral part of decryption. Additionally, for the combined mode, we need to provide two separate keys (encryption and authentication), whereas the AEAD mode generally derives the authentication key internally. Both mentioned integrity methods calculate an *authentication tag* from the final ciphertext (encrypt-then-MAC).

The encryption operation output consists of the encrypted data and the authentication tag. Authentication mode with additional data (AEAD) calculates the authentication tag not only from the input data but also from additional metadata, called additional authentication data (AAD). Table 4 summarizes examples of the encryption modes mentioned in this text.

3.2 Initialization Vectors

IV	Description
plain64	Sector as a 64-bit number (device offset). Used for the XTS mode [4].
ESSIV	Encrypted Salt-Sector IV. The sector number is encrypted using a salt as the key. The salt is derived from the device key with a hash function. Used for the CBC mode [4].
random	IV generated on every sector write from a Random Number Generator. Used for AEAD modes.

Table 5: Initialization vectors.

The Initialization Vector (IV) is a value for an encryption mode that tweaks encryption. In FDE, the IV must always be derived from a sector number (offset from the device start) to prevent malicious sector relocation. The sector number guarantees that the same data written to different sectors produce different ciphertexts. The proper use of IVs and nonces depends on the exact encryption mode and is *critical* for the security

of the entire solution. Table 5 briefly describes the IV types used in our work. For the CBC mode, we must use an IV that an adversary cannot predict; otherwise, the IV value can be used (in combination with a specially formatted plaintext) to create special patterns in the ciphertext (watermarks) [34, 21]. Some encryption modes (such as XTS) solve this problem by encrypting the IV such that they can use a predictable sector number directly.

The IV must always be unique per sector. In some cases, the IV must be a *nonce* (a public value that is never reused). Repeating an IV for different sectors not only opens a possibility to malicious sector relocation but can also violate a security restriction (in the GCM mode repeating a nonce value is fatal [23]).

We define a new *random-IV* that regenerates its value on every write operation by reading it from a system Random Number Generator. The random-IV length should be at least 128 bits to avoid collision. Since stored values are visible to an adversary, the random-IV cannot be used for modes that require unpredictable IVs.

3.3 Error Propagation in Encrypted Sector

Fig. 1: Error propagation in encrypted sector.

With the symmetric encryption in the processing stack, error propagation is amplified. One bit flip causes a random corruption in at least one cipher block (typically 16 bytes) up to the entire sector. This is illustrated in Figure 1.

Such a “random corruption” means that decrypted data are a product of the decryption of a modified ciphertext. By definition, modern ciphers [20] must produce a random-looking output. In other words, a user will see a block full of data garbage after decrypting corrupted data.

For encryption, the change propagation is, in fact, a desirable effect. The ideal situation is that any change in the sector data is propagated to a pseudo-random change to the whole encrypted sector.

4 Metadata Storage Placement

The absence of per-sector metadata (to store integrity protection data) is a well-known problem. Integrity protection requires a length expansion of the processed data and thus needs additional storage [36]. Common sector sizes are 512 and 4096 bytes, and we need an independent metadata per-sector.

4.1 Metadata in Hardware Sector

A reliable way to handle integrity metadata is to do so directly in the device hardware sector. An in-sector integrity data approach appeared in 2003 as the T10 Data Integrity Field (DIF) extension for SCSI devices [24]. This idea was implemented by vendors as T10 Protection Information (T10-PI) where a sector is expanded by 8 bytes (the sector size is 520 bytes) [31]. For COTS devices, the DIF extension is quite rare, expensive and requires a special controller. The fixed provided metadata space is not large enough for cryptographic data integrity protection that requires storing an Initialization Vector to metadata.

4.2 Metadata Stored Separately

Per-sector metadata can be stored in a separate storage space on the same disk or an external fast storage [36]. We can also join several smaller hardware sectors and metadata into one large virtual sector presented to the upper layer. This approach is used in authenticated modes in the FreeBSD GEOM encryption system [25, 14]. The initial GEOM Based Disk Encryption (GBDE) [25] tried to use additional metadata for generated per-sector keys. This design does not provide safe atomic sector updates and also significantly decreases performance

(throughput is only 20-25% of an underlying device) [16]. Some of these problems were fixed by the GELI disk encryption [14]. Here, the upper layer uses a 4096-byte sector, while internally, it splits data into 512-byte native sectors (each contains its own data and metadata). In this case, for every presented sector, 8+1 native sectors are used. This concept ensures that integrity data are handled on the atomic hardware sectors. Although the above filesystem sees the 4096-byte sector as an atomic unit, the device operates with 512-byte units and can possibly interleave multiple writes. Another problem with this design is that the virtual sector is always larger than the native device sector. If the native sector is 4096 bytes, then the presented virtual sector can be larger than the optimal block size for the filesystem above (for some filesystems, the optimal size can be a page size, and thus, it is 4096 bytes in most situations).

Enterprise storage vendors also implement a multiple-sector schema, generally with 8 data sectors + 1 common metadata sector [9].

4.3 Interleaved Metadata Sectors

Metadata [bytes]	Space [%] 512B sector	Space [%] 4096B sector
4 (32 bits)	0.78	0.10
16 (128 bits)	3.03	0.39
32 (256 bits)	5.88	0.78
64 (512 bits)	11.11	1.54

Table 6: Space for per-sector metadata.

Our solution combines the use of a device integrity profile with a metadata per-sector stored independently in dedicated sectors. The presented sector size is configurable. We can use the same size as a native device, but we can also increase the presented sector size (atomicity is then ensured by journaling, as described in Section 4.4). The combined data and metadata are then presented to the block layer as a new virtual device with a specific integrity profile (software-emulated DIF-enabled device).

Fig. 2: Interleaved metadata sectors.

The metadata are stored in special sectors that are interleaved with data sectors. According to the required metadata size, one sector contains a metadata store for several consecutive sectors. An illustration of this layout is provided in Figure 2. Required space examples are illustrated in Table 6. The required size of additional sectors is calculated as follows:

$$TagsPerSector = \lfloor \frac{SectorSize}{TagSize} \rfloor$$

$$Tags_{sectors} = \lceil \frac{Data_{sectors}}{TagsPerSector} \rceil$$

The use of storage is optimal (no wasted space) when the sector size is a multiple of the tag size. Additional metadata space (part of a device used for metadata) is calculated as

$$100 \frac{Tags_{sectors}}{Data_{sectors} + Tags_{sectors}} \%$$

4.4 Recovery on Write Failure

A device must provide atomic updating of both data and metadata. A situation in which one part is written to media while another part failed must not occur. Furthermore, metadata sectors are packed with tags for multiple sectors; thus, a write failure must not cause an integrity validation failure for other sectors.

A metadata-enabled device must implement a data journal that provides a reliable recovery in the case of a power failure. The journal can be switched off if an upper layer provides its own protection of data.

In some specific situations, a journal could provide unintended additional information about the encrypted device (such as last written data offsets or old data content of unfinished sector writes). A journal requires additional storage and decreases performance on write (data are written twice). The journal size is generally just a fraction of the device capacity.

5 Linux Kernel Implementation

We split the solution into *storage* and *cryptographic* parts. The former provides an area for per-sector metadata for commercial off-the-shelf (COTS) storage devices. The latter extends disk encryption by cryptographically sound integrity protection. For key management, we extended the LUKS encryption system [5] to provide support for our integrity protection. Implementation details can be found in the extended technical report [13].

Our approach is based on the existing block layer integrity profile infrastructure [31] and is fully contained in the device-mapper kernel subsystem. The principle of the device-mapper architecture is that separate functions are combined by stacking several devices, each with a specific *target* that implements the needed function. For integrity-enabled encryption, there are three stacked devices: *dm-crypt device* (encrypts and authenticates data), *dm-integrity device* (provides per-sector metadata) and *an underlying block device* (disk or partition). The top-level device is available to a user and can be used directly by an application or an arbitrary filesystem.

5.1 Sector Authenticated Encryption

AAD	DATA	AUTH
authenticated	authenticated + encrypted	TAG
sector	IV	data in/out
		tag

Table 7: AEAD sector authentication request.

The integrity tag is calculated for all sectors, including sectors that are not in use. On the initial activation, the device has to recalculate all integrity tags; otherwise, all

reads fail with an integrity error.

To perform an authenticated encryption operation over a sector, we define the format of a sector authentication request. The additional data (AAD) for our request contain the sector number and the IV.

Such a request detects a sector misplacement and a corrupted IV. The request definition is compatible with the current IEEE 1619 storage standard [2, 17].

5.2 The dm-integrity and the dm-crypt Module

The implemented *dm-integrity* device-mapper target creates a virtual device with arbitrary-sized per-sector metadata over the standard block device and presents it to the system as a device with a specific integrity profile.

The *dm-integrity* target can operate in two modes: as a provider of *per-sector metadata for the upper device* or in *standalone mode*, where it calculates and maintains basic data integrity checksums itself. For reads, it combines the data and metadata sectors and submits them to the upper layer. For writes, it splits data and metadata and submits new I/O requests to the underlying device.

The *dm-crypt* module is a Linux in-kernel implementation of FDE. To support authenticated encryption, we implemented the following parts: *configuration table* format for integrity protection, interface to the underlying I/O *dm-integrity* profile, processing of requests specified in Table 7 and new random-IV generator.

All cryptographic primitives are provided by the kernel cryptographic API. This provides us with all the hardware support and acceleration available. When an integrity check fails, the I/O request is marked as failed and returned to the block layer.

6 Performance

We present two simple benchmarks to illustrate the usability of our solution. The first benchmark measures linear access to the device and the second is a synthetic I/O benchmark with interleaved reads and writes.

We used the mainline kernel 4.13 [6] and cryptsetup tool [5]. The tests ran on a basic installation of the Fedora 26 Linux, were repeated ten times, and the arithmetic mean with the standard deviation (in the format *value* $\pm$ *SD*) is then presented. All cryptographic hardware acceleration modules were loaded.

The hardware configuration represents a typical user configuration that we expect our solution would be used with and was set as follows: Lenovo x240 laptop, Intel Core i7-4600U 2.10 GHz CPU (with AES-NI and SSE3), 8 GB RAM and Toshiba THNSFJ256GCSU 256 GB SSD.

6.1 Linear Access

We ran tests that measured the speed of reads and writes to the entire device. The I/O block size was fixed to 4 kB (presenting a page cache I/O that always produces a page-sized I/O as the minimum).

We ran the tests both with the *dm-integrity* data journal and without the journal (this simulates a situation when a journaling is already present on an upper layer). The measured data are presented in Figure 3.

The measured data show that the overhead of the length-preserving encryption in this scenario is almost negligible. The only visible overhead is for the write operation.

Disabling the data journal has an effect only for write operations. The most visible output is the overhead of additional metadata processing. The overhead

Fig. 3: Device throughput on a solid-state disk.

of cryptographic data integrity processing is visible for all three authenticated modes.

6.2 Random I/O Throughput

We simulated workload with the *fio* [8] (Flexible I/O Tester) utility. We used a *fio* profile simulating mixed reads (70%) and writes (30%) that generates random data access offsets with I/O operations of the 8k block in size. The test uses 16 parallel jobs, asynchronous I/Os and runs for at least 100 seconds.

This test should represent the worst case scenario in our comparison. The measured results are presented in Figure 4. We can see that a journal has a small effect in this scenario; the major visible slowdown is in the additional metadata processing.

7 Conclusions

Our goal was not only to show that the combination of confidentiality and cryptographic data integrity protection is possible at the FDE layer, but also to highlight the need for proper cryptographic data integrity protection in general. We focused on existing COTS devices without any specific hardware requirements. Almost all laptops are currently delivered with an SSD. In this scenario, the performance evaluation shows that our data integrity protection is usable for these systems.

Fig. 4: *fio* simulated load on a solid-state disk.

The *dm-integrity* module provides a generic solution for additional metadata per sector. The price to pay is decreased storage performance and storage capacity, but we believe this is balanced by a zero-cost investment to required hardware. The extension to *dm-crypt* FDE is algorithm-agnostic. The configuration of another encryption mode is just a configuration option.

It appears that ChaCha20 with the Poly1305 authenticator [30] is the authentication encryption algorithm of choice today, yet this can easily change with the CAESAR [3] competition or a new GCM-SIV [23, 28] mode.

Both configurable metadata per-sector and encryption upgrades of algorithms are the major advantages to hardware-based encryption solutions, where any reconfiguration during their lifetime is almost impossible.

All code presented in this work is released under the open-source GPL2 license and has been included in the Linux mainline kernel since the version 4.12 and userspace cryptsetup tool [5] since the version 2.0.0.

7.1 Future Work

Integrity protection is generally used on a higher layer than device sectors [39]. The *dm-integrity* feature could be used later even for another task such as an application-level integrity protection (application could send its own data to be stored on per-sector level) of storing additional Forward Error Correction codes (the storage could then not only detect integrity problems but also fix basic random data corruptions).

Additionally, the same principle can be applied to new storage device types, such as persistent memory (or any byte-addressable persistent storage), where we can easily increase the size of the additional authentication tag (in principle, we can use virtual sector of any size). In this case, the *dm-integrity* layer can be omitted (the atomicity of data and metadata writes is provided by the storage itself), while the *dm-crypt* cryptographic part of the solution remains the same.

Acknowledgments. The authors thank Arno Wagner, John Strunk, Ondrej Mosnáček, Virgil Gligor and Ric Wheeler for valuable comments.

References

1. FIPS Publication 197, The Advanced Encryption Standard (AES) (2001), U.S. DoC/NIST
2. IEEE Standard for Authenticated Encryption With Length Expansion for Storage Devices. IEEE Std 1619.1-2007 (2008). <https://doi.org/10.1109/IEEESTD.2008.4523925>
3. CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness (2016), competitions.cr.yp.to/caesar.html
4. dm-crypt: Linux device-mapper crypto target (2017), gitlab.com/cryptsetup/cryptsetup/wikis/DMCrypt
5. LUKS: Linux Unified Key Setup (2017), gitlab.com/cryptsetup/cryptsetup
6. Linux mainline kernel archive (2018), kernel.org
7. Alendal, G., et al.: got HW crypto? On the (in)security of a Self-Encrypting Drive series. IACR Cryptology ePrint Archive (2015)
8. Axboe, J.: Flexible I/O Tester (2017), github.com/axboe/fio
9. Bairavasundaram, L.N., et al.: An analysis of data corruption in the storage stack. ACM Transactions on Storage **4**(3) (2008)
10. Bellare, M., et al.: Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. J. Cryptol. **21**(4) (2008)
11. Bernstein, D.J.: ChaCha, a variant of Salsa20 (2008), cr.yp.to/chacha/chacha-20080120.pdf
12. Böck, H., et al.: Nonce-Disrespecting Adversaries: Practical Forgery Attacks on GCM in TLS. IACR Cryptology ePrint Archive (2016)
13. Brož, M., et al.: Practical Cryptographic Data Integrity Protection with Full Disk Encryption. Tech. rep. (2018), <https://gitlab.com/cryptsetup/cryptsetup/wikis/integritytechreport>
14. Dawidek, P.J.: FreeBSD GELI encryption system (2011), github.com/freebsd/freebsd/blob/master/sys/geom/eli/g_eli_integrity.c
15. van Dijk, M., et al.: Offline Untrusted Storage with Immediate Detection of Forking and Replay Attacks. In: Proceedings of the 2007 ACM Workshop on Scalable Trusted Computing. STC '07, ACM (2007)
16. Dowdeswell, R.C.: Initial Analysis of GBDE (2003), <http://www.imrryr.org/~elric/cgd/gbde-analysis2.pdf>
17. Dworkin, M.J.: SP 800-38D. Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC (2007), NIST
18. Dworkin, M.J.: SP 800-38E. Recommendation for Block Cipher Modes of Operation: The XTS-AES Mode for Confidentiality on Storage Devices (2010), NIST

19. Ferguson, N.: AES-CBC + Elephant diffuser: A disk encryption algorithm for Windows Vista (2006), Microsoft Corporation
20. Ferguson, N., et al.: *Cryptography Engineering: Design Principles and Practical Applications*. Wiley Publishing (2010)
21. Fruhwirth, C.: New methods in hard disk encryption. Ph.D. thesis, Institute for Computer Languages Theory and Logic Group, Vienna University of Technology (2005)
22. Gjøsteen, K.: Security Notions for Disk Encryption. In: *ESORICS 2005: 10th European Symposium on Research in Computer Security*. Springer Berlin Heidelberg (2005)
23. Gueron, S., et al.: GCM-SIV: Full Nonce Misuse-Resistant Authenticated Encryption at Under One Cycle per Byte. *Cryptology ePrint Archive* (2015)
24. Holt, K.: End-to-End Data Protection Justification (2003), www.t10.org/ftp/t10/document.03/03-224r0.pdf, T10 Technical Committee proposal letter
25. Kamp, P.H.: GBDE: GEOM Based Disk Encryption. In: *Proceedings of the BSD Conference 2003 on BSD Conference*. USENIX (2003)
26. Khati, L., et al.: *Full Disk Encryption: Bridging Theory and Practice*. Springer International Publishing (2017)
27. Krioukov, A., et al.: Parity Lost and Parity Regained. In: *Proceedings of the 6th USENIX Conference on File and Storage Technologies*. USENIX Association (2008)
28. Lindell, Y., et al.: AES-GCM-SIV: Nonce Misuse-Resistant Authenticated Encryption. Internet-Draft draft-irtf-cfrg-gcmsiv-03 (2017)
29. Martin, T., et al.: The 2016 Analysis of Information Remaining on Computer Hard Disks Offered for Sale on the Second Hand Market in the UAE. *Journal of Digital Forensics, Security and Law* 4(4) (2016)
30. Nir, Y., et al.: ChaCha20 and Poly1305 for IETF Protocols. RFC 7539 (2015)
31. Petersen, M.K.: T10 Data Integrity Feature (Logical Block Guarding) (2007), Linux Storage & Filesystem Workshop
32. Riedel, E., et al.: A Framework for Evaluating Storage System Security. In: *Proceedings of the 1st USENIX Conference on File and Storage Technologies*. vol. 2. USENIX Association (2002)
33. Rogaway, P.: Authenticated-encryption with associated-data. In: *Proceedings of the 9th ACM Conference on Computer and Communications Security* (2002)
34. Saarinen, M.J.O.: Encrypted Watermarks and Linux Laptop Security. In: *Workshop on Information Security Applications, Revised Selected Papers* (2005)
35. Satran, J., et al.: Internet Protocol Small Computer System Interface (iSCSI) Cyclic Redundancy Check (CRC)/Checksum Considerations. RFC 7539 (2015)
36. Sivathanu, G., et al.: Ensuring data integrity in storage: Techniques and applications. In: *ACM Workshop on Storage Security and Survivability* (2005)
37. Tischer, M., et al.: Users Really Do Plug in USB Drives They Find. In: *2016 IEEE Symposium on Security and Privacy (SP)* (2016)
38. Törpe, S., et al.: Attacking the BitLocker Boot Process. In: *Proceedings of the 2nd International Conference on Trusted Computing*. Trust '09, Springer-Verlag, Berlin, Heidelberg (2009)
39. Zhang, Y., et al.: End-to-end Data Integrity for File Systems: A ZFS Case Study. In: *Proceedings of the 8th USENIX Conference on File and Storage Technologies*. USENIX Association (2010)