This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCC.2017.2670559, IEEE Transactions on Cloud Computing

Securing Cloud Data under Key Exposure

Ghassan O. Karame, *Member, IEEE*, Claudio Soriente, *Member, IEEE*, Krzysztof Lichota, Srdjan Capkun, *Senior Member, IEEE.*

Abstract—Recent news reveal a powerful attacker which breaks data confidentiality by acquiring cryptographic keys, by means of coercion or backdoors in cryptographic software. Once the encryption key is exposed, the only viable measure to preserve data confidentiality is to limit the attacker's access to the ciphertext. This may be achieved, for example, by spreading ciphertext blocks across servers in multiple administrative domains—thus assuming that the adversary cannot compromise all of them. Nevertheless, if data is encrypted with existing schemes, an adversary equipped with the encryption key, can still compromise a single server and decrypt the ciphertext blocks stored therein. In this paper, we study data confidentiality against an adversary which knows the encryption key and has access to a large fraction of the ciphertext blocks. To this end, we propose Bastion, a novel and efficient scheme that guarantees data confidentiality even if the encryption key is leaked and the adversary has access to almost all ciphertext blocks. We analyze the security of Bastion, and we evaluate its performance by means of a prototype implementation. We also discuss practical insights with respect to the integration of Bastion in commercial dispersed storage systems. Our evaluation results suggest that Bastion is well-suited for integration in existing systems since it incurs less than 5% overhead compared to existing semantically secure encryption modes.

Index Terms—Key exposure, data confidentiality, dispersed storage.

1 Introduction

The world recently witnessed a massive surveillance program aimed at breaking users' privacy. Perpetrators were not hindered by the various security measures deployed within the targeted services [31]. For instance, although these services relied on encryption mechanisms to guarantee data confidentiality, the necessary keying material was acquired by means of backdoors, bribe, or coercion.

If the encryption key is exposed, the only viable means to guarantee confidentiality is to limit the adversary's access to the ciphertext, e.g., by spreading it across multiple administrative domains, in the hope that the adversary cannot compromise all of them. However, even if the data is encrypted and dispersed across different administrative domains, an adversary equipped with the appropriate keying material can compromise a server in one domain and decrypt ciphertext blocks stored therein.

In this paper, we study data confidentiality against an adversary which knows the encryption key and has access to a large fraction of the ciphertext blocks. The adversary can acquire the key either by exploiting flaws or backdoors in the key-generation software [31], or by compromising the devices that store the keys (e.g., at the user-side or in the cloud). As far as we are aware, this adversary invalidates the security of most cryptographic solutions, including those that protect encryption keys by means of secret-sharing (since these keys can be leaked as soon as they are generated).

To counter such an adversary, we propose Bastion, a novel and efficient scheme which ensures that plaintext data cannot be recovered as long as the adversary has access to at most all but *two* ciphertext blocks, even when the encryption key is exposed. Bastion achieves this by combining the use of standard encryption functions with an efficient linear transform. In this sense, Bastion shares similarities with the notion of all-or-nothing transform. An AONT is not an encryption by itself, but can be used as a pre-processing step before encrypting the data with a block cipher. This encryption paradigm—called AON encryption—was mainly intended to slow down brute-force attacks on the encryption key. However, AON encryption can also preserve data confidentiality in case the encryption key is exposed, as long as the adversary has access to at most all but one ciphertext blocks. Existing AON encryption schemes, however, require *at least* two rounds of block cipher encryptions on the data: one pre-processing round to create the AONT, followed by another round for the actual encryption. Notice that these rounds are sequential, and cannot be parallelized. This results in considerable—often unacceptable—overhead to encrypt and decrypt large files. On the other hand, Bastion requires only one round of encryption—which makes it well-suited to be integrated in existing dispersed storage systems.

We evaluate the performance of Bastion in comparison with a number of existing encryption schemes. Our results show that Bastion only incurs a negligible per-

- *G. Karame is affiliated with NEC Laboratories Europe, Heidelberg, 69115 Germany. E-mail: ghassan.karame@neclab.eu*
- *C. Soriente and S. Capkun are affiliated with the Compute Science Department of ETH Zurich, 8092, Switzerland. Email: firstname.lastname@inf.ethz.ch*
- *K. Lichota is affiliated with 9livesdata, Poland. Email: lichota@9livesdata.com*

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCC.2017.2670559, IEEE Transactions on Cloud Computing

formance deterioration (less than 5%) when compared to symmetric encryption schemes, and considerably improves the performance of existing AON encryption schemes [12], [26]. We also discuss practical insights with respect to the possible integration of Bastion in commercial dispersed storage systems. Our contributions in this paper can be summarized as follows:

- We propose Bastion, an efficient scheme which ensures data confidentiality against an adversary that knows the encryption key and has access to a large fraction of the ciphertext blocks.
- We analyze the security of Bastion, and we show that it prevents leakage of any plaintext block as long as the adversary has access to the encryption key and to all but two ciphertext blocks.
- We evaluate the performance of Bastion analytically and empirically in comparison to a number of existing encryption techniques. Our results show that Bastion considerably improves (by more than 50%) the performance of existing AON encryption schemes, and only incurs a negligible overhead when compared to existing semantically secure encryption modes (e.g., the CTR encryption mode).
- We discuss practical insights with respect to the deployment of Bastion within existing storage systems, such as the HYDRAstor grid storage system [13], [23].

The remainder of the paper is organized as follows. In Section 2, we define our notation and building blocks. In Section 4, we describe our model and introduce our scheme, Bastion. In Section 5, we analyze our scheme in comparison with a number of existing encryption primitives. In Section 6, we implement and evaluate the performance of Bastion in realistic settings; we also discuss practical insights with respect to the integration of Bastion within existing dispersed storage systems. In Section 7, we overview related work in the area, and we conclude the paper in Section 8.

2 Preliminaries

We adapt the notation of [12] for our settings. We define a block cipher as a map $F : \{0,1\}^k \times \{0,1\}^l \rightarrow \{0,1\}^l$, for positive k and l. If P_l is the space of all $(2^l)!$ l-bits permutations, then for any $a \in \{0,1\}^k$, we have $F(a,\cdot) \in P_l$. We also write $F_a(x)$ to denote $F(a,x)$. We model F as an ideal block cipher, i.e., a block cipher picked at random from $BC(k,l)$, where $BC(k,l)$ is the space of all block ciphers with parameters k and l. For a given block cipher $F \in BC(k,l)$, we denote $F^{-1} \in BC(k,l)$ as $F^{-1}(a,y)$ or as $F_a^{-1}(y)$, for $a \in \{0,1\}^k$.

2.1 Encryption modes

An encryption mode based on a block cipher F/F^{-1} is given by a triplet of algorithms $\prod = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ where:

- $\mathcal{K}$ The key generation algorithm is a probabilistic algorithm which takes as input a security parameter k and outputs a key $a \in \{0,1\}^k$ that specifies F_a and F_a^{-1}.
- $\mathcal{E}$ The encryption algorithm is a probabilistic algorithm which takes as input a message $x \in \{0,1\}^*$, and uses F_a and F_a^{-1} as oracles to output ciphertext y.
- $\mathcal{D}$ The decryption algorithm is a deterministic algorithm which takes as input a ciphertext y, and uses F_a and F_a^{-1} as oracles to output plaintext $x \in \{0,1\}^*$, or $\perp$ if y is invalid.

For correctness, we require that for any key $a \leftarrow \mathcal{K}(1^k)$, for any message $x \in \{0,1\}^*$, and for any $y \leftarrow \mathcal{E}^{F_a, F_a^{-1}}(x)$, we have $x \leftarrow \mathcal{D}^{F_a, F_a^{-1}}(y)$.

Security is defined through the following chosen-plaintext attack (CPA) game adapted for block ciphers:

$\mathbf{Exp}_{\prod}^{ind}(A,b)$
$F \leftarrow BC(k,l)$
$a \leftarrow \mathcal{K}(1^k)$
$x_0, x_1, state \leftarrow A^{\mathcal{E}^{F_a, F_a^{-1}}}(find)$
$y_b \leftarrow \mathcal{E}^{F_a, F_a^{-1}}(x_b)$
$b' \leftarrow A(guess, y_b, state)$

In the *ind* experiment, the adversary has unrestricted oracle access to $\mathcal{E}^{F_a, F_a^{-1}}$ during the "find" stage. At this point, $\mathcal{A}$ outputs two messages of equal length x_0, x_1, and some $state$ information that are passed as input when the adversary is initialized for the "guess" stage (e.g., $state$ can contain the two messages x_0, x_1). During the "guess" stage, the adversary is given the ciphertext of one message out of x_0, x_1 and must guess which message was actually encrypted. The advantage of the adversary in the *ind* experiment is:

$$\mathrm{Adv}_{\prod}^{ind}(\mathcal{A}) = |\Pr[\mathrm{Exp}_{\prod}^{ind}(A,0) = 1] - \Pr[\mathrm{Exp}_{\prod}^{ind}(A,1) = 1]|$$

Definition 1. An encryption mode $\prod = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ is *ind* secure if for any probabilistic polynomial time (p.p.t.) adversary $\mathcal{A}$, we have $\mathrm{Adv}_{\prod}^{\mathrm{ind}}(\mathcal{A}) \leq \epsilon$, where ϵ is a negligible function in the security parameter.

Remark 1. The *ind* experiment allows the adversary to see the entire (challenge) ciphertext. In a scenario where ciphertext blocks are dispersed across a number of storage servers, this means that the ind-adversary can compromise all storage servers and fetch the data stored therein.

Remark 2. In the *ind* experiment (and in other experiments used in this paper), we adopt the Shannon Model of a block cipher that, in practice, instantiates an independent random permutation for every different key. This model has been used in previous

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCC.2017.2670559, IEEE Transactions on Cloud Computing

related work [3], [12], [17] to disregard the algebraic or cryptanalysis specific to block ciphers and treat them as a black-box transformation.

2.2 All or Nothing Transforms

An All or Nothing Transform (AONT) is an efficiently computable transform that maps sequences of input blocks to sequences of output blocks with the following properties: *(i)* given all output blocks, the transform can be efficiently inverted, and *(ii)* given all but one of the output blocks, it is infeasible to compute any of the original input blocks. The formal syntax of an AONT is given by a pair of p.p.t. algorithms $\prod = (\mathbb{E}, \mathbb{D})$ where:

- $\mathbb{E}$ The encoding algorithm is a probabilistic algorithm which takes as input a message $x \in \{0,1\}^*$, and outputs a pseudo-ciphertext y.
- $\mathbb{D}$ The decoding algorithm is a deterministic algorithm which takes as input a pseudo-ciphertext y, and outputs either a message $x \in \{0,1\}^*$ or $\perp$ to indicate that the input pseudo-ciphertext is invalid.

For correctness, we require that for all $x \in \{0,1\}^*$, and for all $y \leftarrow \mathbb{E}(x)$, we have $x \leftarrow \mathbb{D}(y)$.

The literature comprises a number of security definitions for AONT (e.g., [8], [12], [26]). In this paper, we rely on the definition of [12] which uses the *aont* experiment below. This definition specifies a block length l such that the pseudo-ciphertext y can be written as $y = y[1] \dots y[n]$, where $|y[i]| = l$ and $n \geq 1$.

$\mathbf{Exp}^{aont}_{\prod}(A, b)$
$x, state \leftarrow A(find)$
$y_0 \leftarrow \mathbb{E}(x)$
$y_1 \leftarrow \{0,1\}^{|y_0|}$
$b' \leftarrow A^{Y_b}(guess, state)$

On input j, the oracle Y_b returns $y_b[j]$ and accepts up to $(n-1)$ queries. The *aont* experiment models an adversary which must distinguish between the encoding of a message of its choice and a random string (of the same length), while the adversary is allowed access to all but one encoded blocks. The advantage of $\mathcal{A}$ in the *aont* experiment is given by:

$$\mathrm{Adv}^{aont}_{\prod}(\mathcal{A}) = |\Pr[\mathrm{Exp}^{aont}_{\prod}(A,0) = 1] - \Pr[\mathrm{Exp}^{aont}_{\prod}(A,1) = 1]|$$

Definition 2. An All-or-Nothing Transform $\prod = (\mathbb{E}, \mathbb{D})$ is *aont* secure if for any p.p.t. adversary $\mathcal{A}$, we have $\mathrm{Adv}^{\mathrm{aont}}_{\prod}(\mathcal{A}) \leq \epsilon$, where ϵ is a negligible function in the security parameter.

Known AONTs

Rivest [26] suggested the *package transform* which leverages a block cipher F/F^{-1} and maps m block strings to $n = m + 1$ block strings. The first $n-1$ output blocks are computed by XORing the i-th plaintext block with $F_K(i)$, where K is a random key. The n-th output block is computed XORing K with the encryption of each of the previous output blocks, using a key K_0 that is publicly known. That is, given $x[1] \dots x[m]$, the package transform outputs $y[1] \dots y[n]$, with $n = m + 1$, where:

$$y[i] = x[i] \oplus F_K(i),\ 1 \leq i \leq n-1,$$
$$y[n] = K \bigoplus_{i=1}^{n-1} F_{K_0}(y[i] \oplus i).$$

Desai [12] proposed a faster version where the block cipher round which uses K_0 is skipped and the last output block is set to $y[n] = K \bigoplus_{i=1}^{n-1} y[i]$. Both AONTs are secure according to Definition 2 [12].

Remark 3. Although most proposed AONTs are based on block ciphers [12], [26], an AONT is not an encryption scheme, because there is no secret-key information associated with the transform. Given all the output blocks of the AONT, the input can be recovered without knowledge of any secret.

3 System and Security Model

In this section, we start by detailing the system and security models that we consider in the paper. We then argue that existing security definitions do not capture well the assumption of key exposure, and propose a new security definition that captures this notion.

3.1 System Model

We consider a multi-cloud storage system which can leverage a number of commodity cloud providers (e.g., Amazon, Google) with the goal of distributing trust across different administrative domains. This "cloud of clouds" model is receiving increasing attention nowadays [4], [6], [32] with cloud storage providers such as EMC, IBM, and Microsoft, offering products for multi-cloud systems [15], [16], [29].

In particular, we consider a system of s storage servers $S_1, \dots, S_s$, and a collection of users. We assume that each server appropriately authenticates users. For simplicity and without loss of generality, we focus on the read/write storage abstraction of [21] which exports two operations:

- write(v) This routine splits v into s pieces $\{v_1, \dots, v_s\}$ and sends $\langle v_j \rangle$ to server S_j, for $j \in [1 \dots s]$.
- read$(\cdot)$ The read routine fetches the stored value v from the servers. For each $j \in [1 \dots s]$, piece v_j is downloaded from server S_j and all

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCC.2017.2670559, IEEE Transactions on Cloud Computing

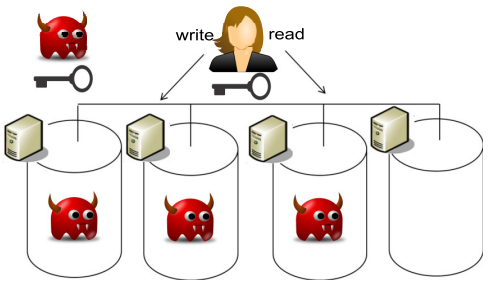

Fig. 1. Our attacker model. We assume an adversary which can acquire all the cryptographic secret material, and can compromise a large fraction (up to all but one) of the storage servers.

pieces are combined into v. We assume that the initial value of the storage is a special value $\perp$, which is not a valid input value for a write operation.

3.2 Adversarial Model

We assume a computationally-bounded adversary $\mathcal{A}$ which *can acquire the long-term cryptographic keys used to encrypt the data*. The adversary may do so either *(i)* by leveraging flaws or backdoors in the key-generation software [31], or *(ii)* by compromising the device that stores the keys (in the cloud or at the user). Since ciphertext blocks are distributed across servers hosted within different domains, we assume that the adversary cannot compromise all storage servers (cf. Figure 1). In particular, we assume that the adversary can compromise all but one of the servers and we model this adversary by giving it access to all but λ ciphertext blocks.

Note that if the adversary also learns the user's credentials to log into the storage servers and downloads all the ciphertext blocks, then no cryptographic mechanism can preserve data confidentiality. We stress that compromising the encryption key does not necessarily imply the compromise of the user's credentials. For example, encryption can occur on a specific-purpose device [10], and the key can be leaked, e.g., by the manufacturer; in this scenario, the user's credentials to access the cloud servers are clearly not compromised.

3.3 $(n-\lambda)$-CAKE Security

Existing security notions for encryption modes capture data confidentiality against an adversary which does not have the encryption key. That is, if the key is leaked, the confidentiality of data is broken.

In this paper we study an adversary that has access to the encryption key but does not have the entire ciphertext. We therefore propose a new security definition that models our scenario.

As introduced above, we allow the adversary to access an encryption/decryption oracle and to "see" all but λ ciphertext blocks. Since confidentiality with $\lambda = 0$ is clearly not achievable[1], we instead seek an encryption mode where $\lambda = 1$. However, having the flexibility of setting $\lambda \geq 1$ allows the design of more efficient schemes while keeping a high degree of security in practical deployments. (See Remark 7.)

We call our security notion $(n-\lambda)$ Ciphertext Access under Key Exposure, or $(n-\lambda)$*CAKE*. Similar to [12], $(n-\lambda)$*CAKE* specifies a block length l such that a ciphertext y can be written as $y = y[1]\dots y[n]$ where $|y[i]| = l$ and $n > 1$.

$$\begin{aligned}&\mathbf{Exp}_{\prod}^{(n-\lambda)CAKE}(A,b)\\ &\quad a \leftarrow \mathcal{K}(1^k)\\ &\quad x_0, x_1, state \leftarrow A^{\mathcal{E}^{F_a,F_a^{-1}}}(find)\\ &\quad y_b \leftarrow \mathcal{E}^{F_a,F_a^{-1}}(x_b)\\ &\quad b' \leftarrow A^{Y_b,\mathcal{E}^{F_a,F_a^{-1}}}(guess, state)\end{aligned}$$

The adversary has unrestricted access to $\mathcal{E}^{F_a,F_a^{-1}}$ in both the "find" and "guess" stages. On input j, the oracle Y_b returns $y_b[j]$ and accepts up to $n-\lambda$ queries. On the one hand, unrestricted oracle access to $\mathcal{E}^{F_a,F_a^{-1}}$ captures the adversary's knowledge of the secret key. On the other hand, the oracle Y_b models the fact that the adversary has access to all but λ ciphertext blocks. This is the case when, for example, each server stores λ ciphertext blocks and the adversary cannot compromise all servers. The advantage of the adversary is defined as:

$$\begin{aligned}Adv_{\prod}^{(n-\lambda)CAKE}(A) = Pr[\mathbf{Exp}_{\prod}^{(n-\lambda)CAKE}(A,1)=1] -\\ Pr[\mathbf{Exp}_{\prod}^{(n-\lambda)CAKE}(A,0)=1]\end{aligned}$$

Definition 3. An encryption mode $\prod = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ is $(n-\lambda)$*CAKE* secure if for any p.p.t. adversary $\mathcal{A}$, we have $\mathrm{Adv}_{\prod}^{(n-\lambda)CAKE}(\mathcal{A}) \leq \epsilon$, where ϵ is a negligible function in the security parameter.

Definition 3 resembles Definition 2 but has two fundamental differences. First, $(n-\lambda)$*CAKE* refers to a keyed scheme and gives the adversary unrestricted access to the encryption/decryption oracles. Second, $(n-\lambda)$*CAKE* relaxes the notion of all-or-nothing and parameterizes the number of ciphertext blocks that are not given to the adversary. As we will show in Section 4.2, this relaxation allows us to design encryption modes that are considerably more efficient than existing modes which offer a comparable level of security.

We stress that $(n-\lambda)$*CAKE* does not consider confidentiality against "traditional" adversaries (i.e., adversaries which do not know the encryption key). Indeed, an *ind*-adversary is not given the encryption key but has access to all ciphertext blocks. That is, the *ind*-adversary can compromise all the s storage servers. An $(n-\lambda)$*CAKE*-adversary is given the encryption key but can access all but λ ciphertext blocks. In practice,

1. Any party with access to all the ciphertext blocks and the encryption key can recover the plaintext.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCC.2017.2670559, IEEE Transactions on Cloud Computing

the $(n-\lambda)$*CAKE*-adversary has the encryption key but can compromise up to $s-1$ storage servers. Therefore, we seek an encryption mode $\prod$ with the following properties:

1) $\prod$ must be *ind* secure against an adversary which does not know the encryption key but has access to all ciphertext blocks (cf. Definition 1), by compromising all storage servers.
2) $\prod$ must be $(n-\lambda)$*CAKE* secure against an adversary which knows the encryption key but has access to $n-\lambda$ ciphertext blocks (cf. Definition 3), since it cannot compromise all storage servers.

Remark 4. Property 2 ensures data confidentiality against the attacker model outlined in Section 3.2. Nevertheless, we must also account for weaker adversaries (i.e., traditional adversaries) that do not know the encryption key but can access the entire ciphertext —hence, *ind* security. Note that if the adversary which has access to the encryption key, can also access all the ciphertext blocks, then no cryptographic mechanism can preserve data confidentiality.

4 Bastion: Security against Key Exposure

In this section, we present our scheme, dubbed Bastion, which ensures that plaintext data cannot be recovered as long as the adversary has access to all but *two* ciphertext blocks—even when the encryption key is exposed. We then analyze the security of Bastion with respect to Definition 1 and Definition 3.

4.1 Overview

Bastion departs from existing AON encryption schemes. Current schemes require a pre-processing round of block cipher encryption for the AONT, followed by another round of block cipher encryption (cf. Figure 2 (a)). Differently, Bastion first encrypts the data with one round of block cipher encryption, and then applies an efficient linear post-processing to the ciphertext (cf. Figure 2 (b)). By doing so, Bastion relaxes the notion of all-or-nothing encryption at the benefit of increased performance (see Figure 2).

More specifically, the first round of Bastion consists of CTR mode encryption with a randomly chosen key K, i.e., $y' = Enc(K, x)$. The output ciphertext y' is then fed to a linear transform which is inspired by the scheme of [28]. Namely, our transform basically computes $y = y' \cdot A$ where A is a square matrix such that: *(i)* all diagonal elements are set to 0, and *(ii)* the remaining off-diagonal elements are set to 1. As we shown later, such a matrix is invertible and has the nice property that $A^{-1} = A$. Moreover, $y = y' \cdot A$ ensures that each input block y'_j will depend on all output blocks y_i except from y_j. This transformation—combined with the fact that the original input blocks have high entropy (due to semantic secure encryption)—result in an *ind*-secure and $(n-2)$*CAKE* secure encryption mode. In the following section, we show how to efficiently compute $y' \cdot A$ by means of bitwise XOR operations.

4.2 Bastion: Protocol Specification

We now detail the specification of Bastion.

On input a security parameter k, the key generation algorithm of Bastion outputs a key $K \in \{0,1\}^k$ for the underlying block-cipher. Bastion leverages block cipher encryption in the CTR mode, which on input a plaintext bitstream x, divides it in blocks $x[1], \ldots, x[m]$, where m is odd[2] such that each block has size l.[3] The set of input blocks is encrypted under key K, resulting in ciphertext $y' = y'[1], \ldots, y'[m+1]$, where $y'[m+1]$ is an initialization vector which is randomly chosen from $\{0,1\}^l$.

Next, Bastion applies a linear transform to y' as follows. Let $n = m+1$ and assume A to be an n-by-n matrix where element $a_{i,j} = 0^l$ if $i = j$ or $a_{i,j} = 1^l$, otherwise.[4] Bastion computes $y = y' \cdot A$, where additions and multiplications are implemented by means of XOR and AND operations, respectively. That is, $y[i] \in y$ is computed as $y[i] = \bigoplus_{j=1}^{j=n} (y'[j] \wedge a_{j,i})$, for $i = 1 \ldots, n$.

Given key K, inverting Bastion entails computing $y' = y \cdot A^{-1}$ and decrypting y' using K. Notice that matrix A is invertible and $A = A^{-1}$. The pseudocode of the encryption and decryption algorithms of Bastion are shown in Algorithms 1 and 2, respectively. Both algorithms use F to denote a generic block cipher (e.g., AES).

In our implementation, we efficiently compute the linear transform using $2n$ XOR operations as follows:

$$t = y'[1] \oplus y'[2] \oplus \cdots \oplus y'[n],$$
$$y[i] = t \oplus y'[i],\ 1 \leq i \leq n.$$

Note that $y'[1] \ldots y'[n]$ (computed up to line 6 in Algorithm 1) are the outputs of the CTR encryption mode, where $y'[n]$ is the initialization vector. Similar to the CTR encryption mode, the final output of Bastion is one block larger than the original input.

4.3 Correctness Analysis

We show that for every $x \in \{0,1\}^{lm}$ where m is odd, and for every $K \in \{0,1\}^l$, we have $x = Dec(K, Enc(K, x))$.

In particular, notice that lines 2-6 of Algorithm 1 and lines 9-12 of Algorithm 2 correspond to the standard CTR encryption and decryption routines, respectively.

2. This requirement is essential for the correctness of the subsequent linear transform on the ciphertext blocks. That is, if m is even, then the transform is not invertible.

3. l is the block size of the particular block cipher used.

4. 0^l and 1^l denote a bitstring of l zeros and a bitstream of l ones, respectively.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCC.2017.2670559, IEEE Transactions on Cloud Computing

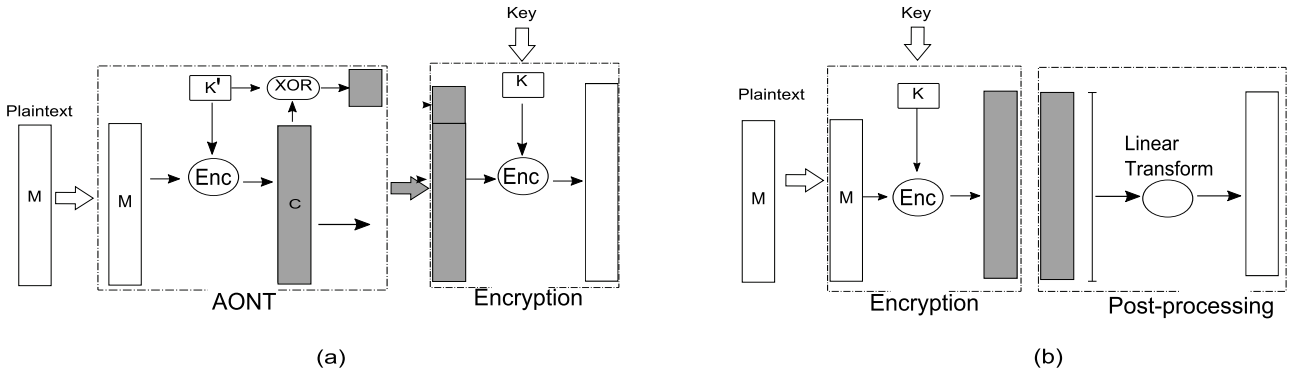

Fig. 2. (a) Current AON encryption schemes require a pre-processing round of block cipher encryption for the AONT, followed by another round of block cipher encryption. (b) On the other hand, Bastion first encrypts the data with one round of block cipher encryption, and then applies an efficient linear post-processing to the ciphertext.

Algorithm 1 Encryption in Bastion.

1: **procedure** $\mathrm{Enc}(K, x = x[1] \dots x[m])$
2: $n = m + 1$
3: $y'[n] \leftarrow \{0,1\}^l$ ▷ y'[n] is the IV for CTR
4: **for** $i = 1 \dots n-1$ **do**
5: $y'[i] = x[i] \oplus F_K(y'[n] + i)$
6: **end for**
7: $t = 0^l$
8: **for** $i = 1 \dots n$ **do**
9: $t = t \oplus y'[i]$
10: **end for**
11: **for** $i = 1 \dots n$ **do**
12: $y[i] = y'[i] \oplus t$
13: **end for**
14: **return** y ▷ $y = y[1] \dots y[n]$
15: **end procedure**

Algorithm 2 Decryption in Bastion.

1: **procedure** $\mathrm{Dec}(K, y = y[1] \dots y[n])$
2: $t = 0^l$
3: **for** $i = 1 \dots n$ **do**
4: $t = t \oplus y[i]$
5: **end for**
6: **for** $i = 1 \dots n$ **do**
7: $y'[i] = y[i] \oplus t$
8: **end for**
9: **for** $i = 1 \dots n-1$ **do**
10: $x[i] = y'[i] \oplus F_K^{-1}(y'[n] + i)$
11: **end for**
12: **return** x ▷ $x = x[1] \dots x[n-1]$
13: **end procedure**

Therefore, we are only left to show that the linear transformation computed in lines 7-14 of Algorithm 1 is correctly reverted in lines 2-8 of Algorithm 2. In other words, we need to show that $t = \bigoplus_{i=1..n} y[i]$ (as computed in the decryption algorithm) matches $t = \bigoplus_{i=1..n} y'[i]$ (as computed in the encryption algorithm).

Recall that t can be computed as follows:

$$
\begin{aligned}
t &= \bigoplus_{i=1..n} y[i] \\
&= \bigoplus_{i=1..n} (y'[i] \oplus t) \\
&= \bigoplus_{i=1..n} \left(y'[i] \oplus \left(\bigoplus_{i=1..n} y'[i] \right) \right) \\
&= \bigoplus_{i=1..n} \left(\bigoplus_{j=1..n, j \neq i} y'[j] \right) \\
&= \bigoplus_{i=1..n} y'[i]
\end{aligned}
$$

Notice that the last step holds because n is even and therefore each $y'[j]$ is XORed for an odd number of times.

Remark 5. We point out that Bastion is not restricted to the CTR encryption mode and can be instantiated with other ind-secure block cipher (and stream ciphers) modes of encryption (e.g., CBC, OFB).

To interface with our cloud storage model described in Section 3.1, we assume that each user encrypts the data using Bastion before invoking the write() routine. More specifically, let $\mathrm{Enc}(K, \cdot), \mathrm{Dec}(K, \cdot)$ denote the encryption and decryption routines of Bastion, respectively. Given encryption key K and a file f, the user computes $v \leftarrow \mathrm{Enc}(K, f)$ and invokes write(v) in order to upload the encrypted file to the cloud. In this setting, key K remains stored at the user's machine. Similarly, to download the file from the cloud, the user invokes read$(\cdot)$ to fetch v and runs $f \leftarrow \mathrm{Dec}(K, v)$ to recover f.

4.4 Security Analysis

In this section, we show that Bastion is $mathrmind$ secure and $(n-2)CAKE$ secure.

Lemma 1. Bastion is ind secure.

Proof 1. Bastion uses an *ind* secure encryption mode to encrypt a message, and then applies a linear

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCC.2017.2670559, IEEE Transactions on Cloud Computing

transform on the ciphertext blocks. It is straightforward to conclude that Bastion is *ind* secure. In other words, a polynomial-time algorithm $\mathcal{A}$ that has non-negligible advantage in breaking the *ind* security of Bastion can be used as a black-box by another polynomial-time algorithm $\mathcal{B}$ to break the *ind* security of the underlying encryption mode. In particular, $\mathcal{B}$ forwards $\mathcal{A}$'s queries to its oracle and applies the linear transformation of Algorithm 1 lines 7-14 to the received ciphertext before forwarding it to $\mathcal{A}$. The same strategy is used when $\mathcal{A}$ outputs two messages at the end of the *find* stage: the two messages are forwarded to $\mathcal{B}$'s oracle; upon receiving the challenge ciphertext, $\mathcal{B}$ applies the linear transformation and forwards it to $\mathcal{A}$. When $\mathcal{A}$ replies with its guess b', $\mathcal{B}$ outputs the same guess. It is easy to see that if $\mathcal{A}$ has non-negligible advantage in guessing correctly which message was encrypted, so does $\mathcal{B}$. Furthermore, the running time of $\mathcal{B}$ is the one of $\mathcal{A}$ plus the time to apply the linear transformation to $\mathcal{A}$'s queries.

Lemma 2. Given any $n-2$ blocks of $y[1]\dots y[n]$ as output by Bastion, it is infeasible to compute any $y'[i]$, for $1 \leq i \leq n$.

Proof 2. Let $y = y[1],\dots,y[n] \leftarrow \mathcal{E}(K, x = x[1]\dots x[m])$. Note that given any $(n-1)$ blocks of y, the adversary can compute one block of y'. In particular, $y'[i] = \bigoplus_{j=1,j\neq i}^{j=n} y[j]$, for any $1 \leq i \leq n$. As it will become clear later, with one block $y'[i]$ and the encryption key, the adversary has non-negligible probability of winning the game of Definition 3.
However, if only $(n-2)$ blocks of y are given, then each of the n blocks of y' can take on any possible values in $\{0,1\}^l$, depending on the two unknown blocks of y. Recall that each block $y'[i]$ is dependent on $(n-1)$ blocks of y and it is pseudo-random as output by the CTR encryption mode. Therefore, given any $(n-2)$ blocks of y, then $y'[i]$ could take any of the 2^l possibilities, for $1 \leq i \leq n$.

Lemma 3. Bastion is $(n-2)$*CAKE* secure.

Proof 3. The security proof of Bastion resembles the standard security proof of the CTR encryption mode and relies on the existence of pseudo-random permutations. In particular, given a polynomial-type algorithm $\mathcal{A}$ which has non-negligible advantage in the $(n-\lambda)$*CAKE* experiment with $\lambda = 2$, we can construct a polynomial-time algorithm $\mathcal{B}$ which has non-negligible advantage in distinguishing between a true random permutation and a pseudo-random permutation.
$\mathcal{B}$ has access to oracle $\mathcal{O}$ and uses it to answer the encryption and decryption queries issued by $\mathcal{A}$. In particular, $\mathcal{A}$'s queries are answered as follows:

- *Decryption query for* $y[1]\dots y[n]$
  1) Compute $t = y[1] \oplus \dots \oplus y[n]$
  2) Compute $y'[i] = y[i] \oplus t$, for $1 \leq i \leq n$
  3) Compute $x[i] = y'[i] \oplus O(y'[n]+i)$, for $1 \leq i \leq n-1$
  4) Return $x[1]\dots x[n-1]$
- *Encryption query for* $x[1]\dots x[n-1]$
  1) Pick random $y'[n] \in \{0,1\}^l$
  2) Compute $y'[i] = x[i] \oplus O(y'[n]+i)$, for $1 \leq i \leq n-1$
  3) Compute $t = y'[1] \oplus \dots \oplus y'[n]$
  4) Compute $y[i] = y'[i] \oplus t$, for $1 \leq i \leq n$
  5) Return $y[1]\dots y[n]$

When $\mathcal{A}$ outputs two messages $x_1[1]\dots x_1[n-1]$ and $x_2[1]\dots x_2[n-1]$, $\mathcal{B}$ picks $b \in \{0,1\}$ at random and does the following:

1) Pick random $y'_b[n] \in \{0,1\}^l$
2) Compute $y'_b[i] = x_b[i] \oplus O(y'_b[n], i)$, for $1 \leq i \leq n-1$
3) Compute $t = y'_b[1] \oplus \dots \oplus y'_b[n]$
4) Compute $y_b[i] = y'_b[i] \oplus t$, for $1 \leq i \leq n$

At this point, $\mathcal{A}$ selects $(n-2)$ indexes $i_1,\dots i_{n-2}$ and $\mathcal{B}$ returns the corresponding $y_b[i_1],\dots,y_b[i_{n-2}]$. Encryption and decryption queries are answered as above. When $\mathcal{A}$ outputs its answer b', $\mathcal{B}$ outputs 1 if $b = b'$, and 0 otherwise. It is straightforward to see that if $\mathcal{A}$ has advantage larger than negligible to guess b, then $\mathcal{B}$ has advantage larger than negligible to distinguish a true random permutation from a pseudorandom one. Furthermore, the number of queries issued by $\mathcal{B}$ to its oracle amounts to the number of encryption and decryption queries issued by $\mathcal{A}$. Note that by Lemma 2, during the guess stage, $\mathcal{A}$ cannot issue a decryption query on the challenge ciphertext since with only $(n-2)$ blocks, finding the remaining blocks is infeasible.

Remark 6. Bastion is not $(n-1)$*CAKE* secure. As shown in the proof of Lemma 2, the adversary can recover one block of y' given any $(n-1)$ blocks of y. If the adversary recovers $y'[n]$ that is used as an IV in the CTR encryption mode, the adversary can easily win the $(n-1)$*CAKE* game. Recall that our security definition allows the adversary to learn the encryption key.

Remark 7. Bastion is $(n-2)$*CAKE* secure according to Definition 3. However, in a practical deployment, we expect that each file spans several thousands blocks [5]. When those blocks are evenly spread across servers, each server will store a larger number of blocks. Therefore, an $(n-2)$*CAKE* secure scheme such as Bastion clearly preserves data confidentiality unless *all* servers are compromised.

5. For example, a 10MB file encrypted using AES has more than 600K blocks.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCC.2017.2670559, IEEE Transactions on Cloud Computing

TABLE 1
Comparison between Bastion and existing constructs. We assume a plaintext of $m = n - 1$ blocks. Since all schemes are symmetric, we only show the computation overhead for the encryption/encoding routine in the column "Computation" ("b.c." is the number of block cipher operations; "XOR" is the number of XOR operations.

	Computation	**Storage** (blocks)	**Security**
CTR Encryption	$n-1$ b.c. $n-1$ XOR	n	1*CAKE* ind-secure
Rivest AONT [26]	$2(n-1)$ b.c. $3(n-1)$ XOR	n	N/A ind-INsecure
Desai AONT [12]	$n-1$ b.c. $2(n-1)$ XOR	n	N/A ind-INsecure
Rivest AON Encryption [26]	$3n-2$ b.c. $3(n-1)$ XOR	n	$(n-1)$*CAKE* ind-secure
Desai AON Encryption [12]	$2n-1$ b.c. $2(n-1)$ XOR	n	$(n-1)$*CAKE* ind-secure
Encrypt-then -secret-share	$n-1$ b.c. $2n-1$ XOR	n^2	$(n-1)$*CAKE* ind-INsecure⋆
Bastion	$n-1$ b.c. $3n-1$ XOR	n	$(n-2)$*CAKE* ind-secure

⋆ Recall that an *ind*-adversary can access all storage servers to fetch all ciphertext blocks. Therefore, the adversary can also fetch all the key shares and compute the encryption key.

5 Comparison to Existing Schemes

In what follows, we briefly overview several encryption modes and argue about their security (according to Definitions 1 and 3) and performance when compared to Bastion.

CPA-encryption modes

Traditional CPA-encryption modes, such as the CTR mode, provide *ind* security but are only 1*CAKE* secure. That is, an adversary equipped with the encryption key must only fetch two ciphertext blocks to break data confidentiality.[6]

CPA-encryption and secret-sharing

Another option is to rely on the combination of CPA secure encryption modes and secret-sharing.

If the file f is encrypted and then shared with an n-out-of-n secret-sharing scheme (denoted as "encrypt-then-secret-share" in the following), then the construction is clearly $(n-1)$*CAKE* secure and is also *ind* secure. However, secret-sharing the ciphertext comes at considerable storage costs; for example, each share would be as large as the file f using a perfect secret sharing scheme—which makes it impractical for storing large files.

Secret-sharing the encryption key and dispersing its shares across the storage servers alongside the ciphertext is not secure against an *ind*-adversary. Indeed, if the adversary can access all the storage servers and download all ciphertext blocks, the adversary may as well download all key shares and compute the encryption key.

6. We assume that the CTR encryption routine starts with a random IV that is incremented at every block encryption.

AON encryption

Recall that an AONT is not an encryption scheme and does not require the decryptor to have any secret key. That is, an AONT is not secure against an *ind*-adversary which can access all the ciphertext blocks. One alternative is to combine the use of AONT with standard encryption. Rivest [26] suggests to pre-process a message with an AONT and then encrypt its output with an encryption mode. This paradigm is referred to in the literature as AON encryption and provides $(n-1)$*CAKE* security. Existing AON encryption schemes require at least two rounds of block cipher encryption with two different keys [12], [26]. At least one round is required for the actual AONT that embeds the first encryption key in the pseudo-ciphertext (cf. Section 2). An additional round uses another encryption key that is kept secret to guarantee CPA-security. However, two encryption rounds constitute a considerable overhead when encrypting and decrypting large files. In Appendix A, we describe possible ways of modifying the AONTs of [26] and [12] to achieve *ind* security and $(n-1)$*CAKE* security without adding another round of block cipher encryption, and we discuss their shortcomings.

Clearly, these solutions are either not satisfactory in terms of security or incur a large overhead when compared to Bastion and may not be suitable to store large files in a multi-cloud storage system.

5.1 Performance Comparison

Table 1 compares the performance of Bastion with the encryption schemes considered so far, in terms of computation, storage, and security.

Given a plaintext of m blocks, the CTR encryption mode outputs $n = m + 1$ ciphertext blocks, computed with $(n-1)$ block cipher operations and $(n-1)$ XOR

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCC.2017.2670559, IEEE Transactions on Cloud Computing

operations. The CTR encryption mode is *ind* secure but only 1*CAKE* secure.

Rivest AONT outputs a pseudo-ciphertext of $n = m+1$ blocks using $2(n-1)$ block cipher operations and $3(n-1)$ XOR operations. Desai AONT outputs the same number of blocks but requires only $(n-1)$ block cipher operations and $2(n-1)$ XOR operations. Both Rivest AONT and Desai AONT are, however, not *ind* secure since the encryption key used to compute the AONT output is embedded in the output itself. Encrypting the output of Rivest AONT or Desai AONT with a standard encryption mode (both [12] and [26] use the ECB encryption mode), requires additional n block cipher operations, and yields an AON encryption that is *ind* secure[7] and $(n-1)$*CAKE* secure. Encrypt-then-secret-share (cf. Section 4.4) is *ind* secure and $(n-1)$*CAKE* secure. It requires $(n-1)$ block cipher operations and n XOR operations if additive secret sharing is used. However secret-sharing encryption results in a prohibitively large storage overhead of n^2 blocks.

Bastion also outputs $n = m+1$ ciphertext blocks. It achieves *ind* security and $(n-2)$*CAKE* security with only $(n-1)$ block cipher operations and $(3n-1)$ XOR operations.[8]

We conclude that Bastion achieves a solid tradeoff between the computational overhead of existing AON encryption modes and the exponential storage overhead of secret-sharing techniques, while offering a comparable level of security. In Section 6, we confirm the superior performance of Bastion by means of implementation.

6 Implementation and Evaluation

In this section, we describe and evaluate a prototype implementation modeling a read-write storage system based on Bastion. We also discuss insights with respect to the integration of Bastion within existing dispersed storage systems.

6.1 Implementation Setup

Our prototype, implemented in C++, emulates the read-write storage model of Section 3.1. We instantiate Bastion with the CTR encryption mode (cf. Figure 1) using both AES128 and Rijndael256, implemented using the libmcrypt.so. 4.4.7 library. Since this library does not natively support the CTR encryption mode, we use it for the generation of the CTR keystream, which is later XORed with the plaintext.

We compare Bastion with the AON encryption schemes of Rivest [26] and Desai [12]. For baseline comparison, we include in our evaluation the CTR encryption mode and the AONTs due to Rivest [26] and Desai [12], which are used in existing dispersed storage systems, e.g., Cleversafe [25]. We do not evaluate the performance of secret-sharing the data because of its prohibitively large storage overhead (squared in the number of input blocks). We evaluate our implementations on an Intel(R) Xeon(R) CPU E5-2470 running at 2.30GHz. Note that the processor clock frequency might have been higher during the evaluation due to the TurboBoost technology of the CPU. In our evaluation, we abstract away the effects of network delays and congestion, and we only assess the processing performance of the encryption for the considered schemes. This is a reasonable assumption since all schemes are length-preserving (plus an additional block of l bits), and are therefore likely to exhibit the same network performance. Moreover, we only measure the performance incurred during encryption/encoding, since all schemes are symmetric, and therefore the decryption/decoding performance is comparable to that of the encryption/encoding process.

We measure the peak throughput and the latency exhibited by our implementations w.r.t. various file/block sizes. For each data point, we report the average of 30 runs. Due to their small widths, we do not show the corresponding 95% confidence intervals.

6.2 Evaluation Results

Our evaluation results are reported in Figure 3 and Figure 4. Both figures show that Bastion considerably improves (by more than 50%) the performance of existing $(n-1)$*CAKE* encryption schemes and only incurs a negligible overhead when compared to existing semantically secure encryption modes (e.g., the CTR encryption mode) that are only 1*CAKE* secure.

In Figure 3, we show the peak throughput achieved by the CTR encryption mode, Bastion, Desai AONT/AON, and Rivest AONT/AON schemes. The peak throughput achieved by Bastion reaches almost 72 MB/s and is only 1% lower than the one exhibited by the CTR encryption mode. When compared with existing $(n-1)$*CAKE* secure schemes, such as Desai AON encryption and Rivest AON encryption, our results show that the peak throughput of Bastion is almost twice as large as that of Desai AON encryption, and more than three times larger than the peak throughput of Rivest AON encryption.

We also evaluate the performance of Bastion, with respect to different block sizes of the underlying block cipher. Our results show that—irrespective of the block size—Bastion only incurs a negligible performance deterioration in peak throughput when compared to the CTR encryption mode. Figures 4(a) and 4(b) show the latency (in ms) incurred by the encryption/encoding routines for different file sizes. The latency of Bastion is comparable to that of the CTR encryption mode—for both AES128 and Rijandael256—and results in a considerable improvement over existing AON encryption schemes (more than 50% gain in latency).

7. Security according to Definition 1 is achieved because the key used to create the AONT is always random, even if the key used to add the outer layer of encryption is fixed.

8. Bastion requires $(n-1)$ XOR operations for the CTR encryption and $2n$ XOR operations for the linear transform.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCC.2017.2670559, IEEE Transactions on Cloud Computing

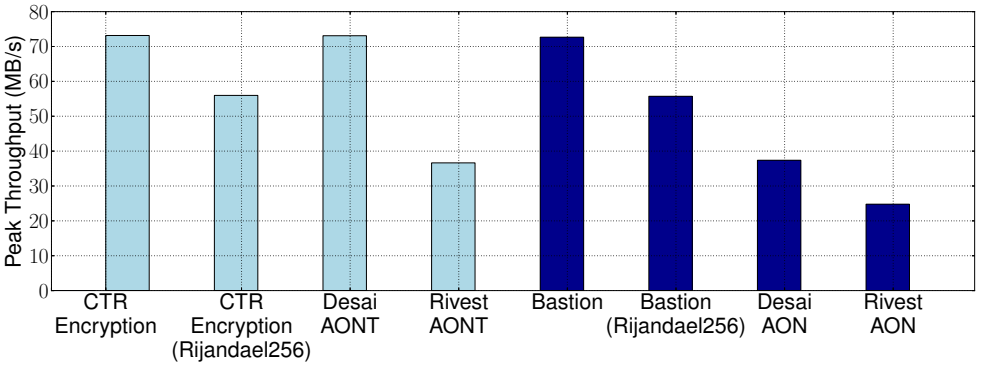

Fig. 3. Peak throughput comparison. Unless otherwise specified, the underlying block cipher is AES128. Each data point is averaged over 30 runs. Histograms in dark blue depict encryption modes which offer comparable security to Bastion. Light blue histograms refer to encryption/encoding modes where individual ciphertext blocks can be inverted when the key is exposed.

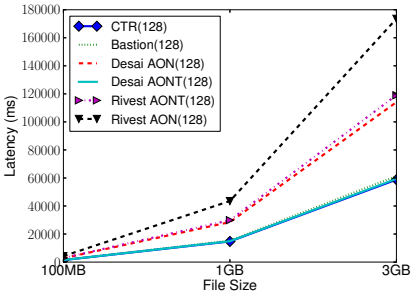

(a) Latency of encryption/encoding for different file sizes.

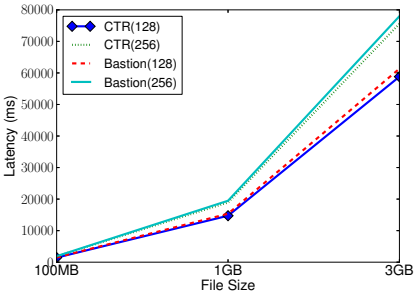

(b) Latency of encryption/encoding for different block sizes of the underlying block cipher.

Fig. 4. Performance evaluation of Bastion. Each data point in is averaged over 30 runs. Unless otherwise specified, the underlying block cipher is AES-128. CTR(256) and Bastion(256) denote the CTR encryption mode and Bastion encryption routine, respectively, instantiated with Rijandael256.

6.3 Deployment within HYDRAstor

Recall that Bastion preserves data confidentiality against an adversary that has the encryption key as long as the adversary does not have access to two ciphertext blocks. In a multi-cloud storage system, if each server stores at least two ciphertext blocks, then Bastion clearly preserves data confidentiality unless *all* servers are compromised.

In scenarios where servers can be faulty, Bastion can be combined with information dispersal algorithms (e.g., [24]) to provide data confidentiality and fault tolerance. Recall that information dispersal algorithms (IDA), parameterized with t_1, t_2 (where $t_1 \leq t_2$), encode data into t_2 symbols such that the original data can be recovered from any t_1 encoded symbols. In our multi-cloud storage system (cf. Section 3.1), the ciphertext output by Bastion is then fed to the IDA encoding routine, with symbols of size l bits, and with parameters $t_2 \geq 2s$, $t_1 < t_2$, where s is the number of available servers. Since the output of the IDA is equally spread across the s servers, by setting $t_2 \geq 2s$, we ensure that each server stores at least two ciphertext blocks worth of data. Finally, the encoded symbols are input to the write() routine that distributes symbols evenly to each of the storage servers. Recovering f via the read() routine entails fetching t_1 encoded symbols from the servers and decoding them via the IDA decoding routine. The resulting ciphertext can be decrypted using Bastion to recover file f. By doing so, data confidentiality is preserved even if the key is exposed unless $t = \frac{st_1}{t_2}$ servers are compromised. Furthermore, data availability is guaranteed in spite of $(s - t)$ server failures.

HYDRAstor

We now discuss the integration of a prototype implementation of Bastion within the HYDRAstor grid storage system [13], [23]. HYDRAstor is a commercial secondary storage solution for enterprises, which consists of a back-end architectured as a grid of storage nodes built around a distributed hash table. HYDRAstor tolerates multiple disk, node and network failures, rebuilds the data automatically after failures, and informs users about recoverability of the deposited data [13]. The reliability and availability of the stored data can be dynamically adjusted by the clients with each write operation, as the back-end supports multiple data resiliency classes [13].

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCC.2017.2670559, IEEE Transactions on Cloud Computing

HYDRAstor distributes written data to multiple disks using the distributed resilient data technology (DRD); the combination of Bastion with DRD ensures that an adversary which has the encryption key and compromises a subset of the disks (i.e., determined by the reconstruction threshold), cannot acquire any meaningful information about the data stored on the disk. To better assess the performance impact of Bastion in HYDRAstor, we evaluated the performance of Bastion in the newest generation HYDRAstor HS8-4000 series system, which uses CPUs with accelerated AES encryption (i.e., the AESNI instruction set). In our experiments, all written data was unique to remove the effect of data deduplication. Results show that the write bandwidth was not affected by the integration of Bastion. The read bandwidth decreased only by 3%. In both read and write operations, the CPU utilization in the system only increased marginally. These experiments clearly suggest that Bastion can be integrated in existing commercial storage systems to strengthen the security of these systems under key exposure, without affecting performance.

7 Related Work

To the best of our knowledge, this is the first work that addresses the problem of securing data stored in multi-cloud storage systems when the cryptographic material is exposed. In the following, we survey relevant related work in the areas of deniable encryption, information dispersal, all-or-nothing transformations, secret-sharing techniques, and leakage-resilient cryptography.

Deniable Encryption

Our work shares similarities with the notion of "shared-key deniable encryption" [9], [14], [18]. An encryption scheme is "deniable" if—when coerced to reveal the encryption key—the legitimate owner reveals "fake keys" thus forcing the ciphertext to "look like" the encryption of a plaintext different from the original one—hence keeping the original plaintext private. Deniable encryption therefore aims to deceive an adversary which does not know the "original" encryption key but, e.g., can only acquire "fake" keys. Our security definition models an adversary that has access to the real keying material.

Information Dispersal

Information dispersal based on erasure codes [30] has been proven as an effective tool to provide reliability in a number of cloud-based storage systems [1], [2], [20], [33]. Erasure codes enable users to distribute their data on a number of servers and recover it despite some servers failures.

Ramp schemes [7] constitute a trade-off between the security guarantees of secret sharing and the efficiency of information dispersal algorithms. A ramp scheme achieves higher "code rates" than secret sharing and features two thresholds t_1, t_2. At least t_2 shares are required to reconstruct the secret and less than t_1 shares provide no information about the secret; a number of shares between t_1 and t_2 leak "some" information.

All or Nothing Transformations

All-or-nothing transformations (AONTs) were first introduced in [26] and later studied in [8], [12]. The majority of AONTs leverage a secret key that is embedded in the output blocks. Once all output blocks are available, the key can be recovered and single blocks can be inverted. AONT, therefore, is not an encryption scheme and does not require the decryptor to have any key material. Resch et al. [25] combine AONT and information dispersal to provide both fault-tolerance and data secrecy, in the context of distributed storage systems. In [25], however, an adversary which knows the encryption key can decrypt data stored on single servers.

Secret Sharing

Secret sharing schemes [5] allow a dealer to distribute a secret among a number of shareholders, such that only authorized subsets of shareholders can reconstruct the secret. In threshold secret sharing schemes [11], [27], the dealer defines a threshold t and each set of shareholders of cardinality equal to or greater than t is authorized to reconstruct the secret. Secret sharing guarantees security against a non-authorized subset of shareholders; however, they incur a high computation/storage cost, which makes them impractical for sharing large files. Rabin [24] proposed an information dispersal algorithm with smaller overhead than the one of [27], however the proposal in [24] does not provide any security guarantees when a small number of shares (less than the reconstruction threshold) are available. Krawczyk [19] proposed to combine both Shamir's [27] and Rabin's [24] approaches; in [19] a file is first encrypted using AES and then dispersed using the scheme in [24], while the encryption key is shared using the scheme in [27]. In Krawczyk's scheme, individual ciphertext blocks encrypted with AES can be decrypted once the key is exposed.

Leakage-resilient Cryptography

Leakage-resilient cryptography aims at designing cryptographic primitives that can resist an adversary which learns partial information about the secret state of a system, e.g., through side-channels [22]. Different models allow to reason about the "leaks" of real implementations of cryptographic primitives [22]. All of these models, however, limit in some way the knowledge of the secret state of a system by the adversary. In contrast, the adversary is given all the secret material in our model.

8 Conclusion

In this paper, we addressed the problem of securing data outsourced to the cloud against an adversary which has access to the encryption key. For that purpose, we introduced a novel security definition that captures data confidentiality against the new adversary.

We then proposed Bastion, a scheme which ensures the confidentiality of encrypted data even when the adversary has the encryption key, and all but *two* ciphertext blocks. Bastion is most suitable for settings where the ciphertext blocks are stored in multi-cloud storage systems. In these settings, the adversary would need to acquire the encryption key, and to compromise *all* servers, in order to recover any single block of plaintext.

We analyzed the security of Bastion and evaluated its performance in realistic settings. Bastion considerably improves (by more than 50%) the performance of existing primitives which offer comparable security under key exposure, and only incurs a negligible overhead (less than 5%) when compared to existing semantically secure encryption modes (e.g., the CTR encryption mode). Finally, we showed how Bastion can be practically integrated within existing dispersed storage systems.

References

[1] M. Abd-El-Malek, G. R. Ganger, G. R. Goodson, M. K. Reiter, and J. J. Wylie, "Fault-Scalable Byzantine Fault-Tolerant Services," in *ACM Symposium on Operating Systems Principles (SOSP)*, 2005, pp. 59–74.

[2] M. K. Aguilera, R. Janakiraman, and L. Xu, "Using Erasure Codes Efficiently for Storage in a Distributed System," in *International Conference on Dependable Systems and Networks (DSN)*, 2005, pp. 336–345.

[3] W. Aiello, M. Bellare, G. D. Crescenzo, and R. Venkatesan, "Security amplification by composition: The case of doubly-iterated, ideal ciphers," in *Advances in Cryptology (CRYPTO)*, 1998, pp. 390–407.

[4] C. Basescu, C. Cachin, I. Eyal, R. Haas, and M. Vukolic, "Robust Data Sharing with Key-value Stores," in *ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC)*, 2011, pp. 221–222.

[5] A. Beimel, "Secret-sharing schemes: A survey," in *International Workshop on Coding and Cryptology (IWCC)*, 2011, pp. 11–46.

[6] A. Bessani, M. Correia, B. Quaresma, F. André, and P. Sousa, "DepSky: Dependable and Secure Storage in a Cloud-of-clouds," in *Sixth Conference on Computer Systems (EuroSys)*, 2011, pp. 31–46.

[7] G. R. Blakley and C. Meadows, "Security of ramp schemes," in *Advances in Cryptology (CRYPTO)*, 1984, pp. 242–268.

[8] V. Boyko, "On the Security Properties of OAEP as an All-or-nothing Transform," in *Advances in Cryptology (CRYPTO)*, 1999, pp. 503–518.

[9] R. Canetti, C. Dwork, M. Naor, and R. Ostrovsky, "Deniable Encryption," in *Proceedings of CRYPTO*, 1997.

[10] Cavalry, "Encryption Engine Dongle," http://www.cavalrystorage.com/en2010.aspx/.

[11] C. Charnes, J. Pieprzyk, and R. Safavi-Naini, "Conditionally secure secret sharing schemes with disenrollment capability," in *ACM Conference on Computer and Communications Security (CCS)*, 1994, pp. 89–95.

[12] A. Desai, "The security of all-or-nothing encryption: Protecting against exhaustive key search," in *Advances in Cryptology (CRYPTO)*, 2000, pp. 359–375.

[13] C. Dubnicki, L. Gryz, L. Heldt, M. Kaczmarczyk, W. Kilian, P. Strzelczak, J. Szczepkowski, C. Ungureanu, and M. Welnicki, "HYDRAstor: a Scalable Secondary Storage," in *USENIX Conference on File and Storage Technologies (FAST)*, 2009, pp. 197–210.

[14] M. Dürmuth and D. M. Freeman, "Deniable encryption with negligible detection probability: An interactive construction," in *EUROCRYPT*, 2011, pp. 610–626.

[15] EMC, "Transform to a Hybrid Cloud," http://www.emc.com/campaign/global/hybridcloud/index.htm.

[16] IBM, "IBM Hybrid Cloud Solution," http://www-01.ibm.com/software/tivoli/products/hybrid-cloud/.

[17] J. Kilian and P. Rogaway, "How to protect DES against exhaustive key search," in *Advances in Cryptology (CRYPTO)*, 1996, pp. 252–267.

[18] M. Klonowski, P. Kubiak, and M. Kutylowski, "Practical Deniable Encryption," in *Theory and Practice of Computer Science (SOFSEM)*, 2008, pp. 599–609.

[19] H. Krawczyk, "Secret Sharing Made Short," in *Advances in Cryptology (CRYPTO)*, 1993, pp. 136–146.

[20] J. Kubiatowicz, D. Bindel, Y. Chen, S. E. Czerwinski, P. R. Eaton, D. Geels, R. Gummadi, S. C. Rhea, H. Weatherspoon, W. Weimer, C. Wells, and B. Y. Zhao, "OceanStore: An Architecture for Global-Scale Persistent Storage," in *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2000, pp. 190–201.

[21] L. Lamport, "On interprocess communication," 1985.

[22] S. Micali and L. Reyzin, "Physically observable cryptography (extended abstract)," in *Theory of Cryptography Conference (TCC)*, 2004, pp. 278–296.

[23] NEC Corp., "HYDRAstor Grid Storage," http://www.hydrastor.com.

[24] M. O. Rabin, "Efficient dispersal of information for security, load balancing, and fault tolerance," *J. ACM*, vol. 36, no. 2, pp. 335–348, 1989.

[25] J. K. Resch and J. S. Plank, "AONT-RS: Blending Security and Performance in Dispersed Storage Systems," in *USENIX Conference on File and Storage Technologies (FAST)*, 2011, pp. 191–202.

[26] R. L. Rivest, "All-or-Nothing Encryption and the Package Transform," in *International Workshop on Fast Software Encryption (FSE)*, 1997, pp. 210–218.

[27] A. Shamir, "How to Share a Secret?" in *Communications of the ACM*, 1979, pp. 612–613.

[28] D. R. Stinson, "Something About All or Nothing (Transforms)," in *Designs, Codes and Cryptography*, 2001, pp. 133–138.

[29] StorSimple, "Cloud Storage," http://www.storsimple.com/.

[30] J. H. van Lint, *Introduction to Coding Theory*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 1982.

[31] Wikipedia, "Edward Snowden," http://en.wikipedia.org/wiki/Edward_Snowden#Disclosure.

[32] Z. Wu, M. Butkiewicz, D. Perkins, E. Katz-Bassett, and H. V. Madhyastha, "SPANStore: Cost-effective Geo-replicated Storage Spanning Multiple Cloud Services," in *ACM Symposium on Operating Systems Principles (SOSP)*, 2013, pp. 292–308.

[33] H. Xia and A. A. Chien, "RobuSTore: a Distributed Storage Architecture with Robust and High Performance," in *ACM/IEEE Conference on High Performance Networking and Computing (SC)*, 2007, p. 44.

Appendix A
Enhancing known AONTs

In what follows, we discuss other means to transform the AONT by Rivest [26] and Desai [12] into an encryption mode that is *ind* secure and $(n-1)$*CAKE* secure, without adding another round of encryption.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCC.2017.2670559, IEEE Transactions on Cloud Computing

A.1 Rivest

Given an input message $x[1] \dots x[m]$, the package transform proposed by Rivest [26] outputs $y[1] \dots y[n]$, with $n = m + 1$, where:

$$y[i] = x[i] \oplus F_K(i),\ 1 \leq i \leq n-1,$$
$$y[n] = K \bigoplus_{i=1}^{n-1} F_{K_0}(y[i] \oplus i).$$

Rivest suggests to choose K uniformly at random for each input message, and to rely on a publicly known K_0. However, it is easy to show that if K_0 is kept secret, then the transform is both *ind* secure, and $(n-1)$*CAKE* secure. This is achieved using $2n-1$ block cipher encryptions and $3(n-1)$ XOR operations. Bastion, on the other hand, only requires $n-1$ block cipher encryptions and $3(n-1)$ XOR operations (cf. Table 1).

A.2 Desai

Desai [12] proposed a faster AONT in which the block cipher round which uses K_0 is skipped and the last output block is set to $y[n] = K \bigoplus_{i=1}^{n-1} y[i]$. One way to achieve both *ind* security and $(n-1)$*CAKE* security is to additionally XOR $y[n]$ with a long-term secret key K_0. As a result, single ciphertext blocks can only be decrypted with key K, and recovery of this key requires all output blocks and knowledge of key K_0. While this proposal is efficient (i.e., it only requires one round of encryption), if a single short-term key K used to encrypt a message is leaked, then the adversary can recover the long-term key K_0. Once K_0 is known, the adversary can decrypt any other message.

Another alternative would be to encrypt one output block (not necessarily the last block, see before) of Desai's AONT with K_0. Although this approach achieves $(n-1)$*CAKE* security, it can only achieve *ind* security if the underlying AONT is *ind* secure.

Ghassan Karame Ghassan Karame is a Senior Researcher at NEC Laboratories Europe. He received his Masters of Science in Information Networking from Carnegie Mellon University (CMU) in December 2006, and his PhD degree in Computer Science from ETH Zurich, Switzerland, in 2011. Between 2011 and 2012, he worked as a postdoctoral researcher in the Institute of Information Security of ETH Zurich. He is a member of the IEEE and of the ACM.

Claudio Soriente Claudio Soriente is a researcher at Telefonica Research and Development. Before his current appointment, he was at the Swiss Federal Institute of Technology, Zuerich, Switzerland and at the Polytechnic University of Madrid, Spain. He received the Ph.D. degree from the University of California, Irvine, CA, USA. His research interests include network and wireless security, privacy, and applied cryptography.

Krzysztof Lichota Krzysztof Lichota is a Senior Technical Expert in 9LivesData LLC. He received his Master of Sciences in Computer Sciences from Warsaw University in 2002. He is working on NEC HydraStor - highly scalable, high performance backup deduplication system and he is co-author of several papers and inventions related to storage..

Srdjan Capkun Srdjan Capkun is an Associate Professor in the Department of Computer Science, ETH Zurich and Director of the Zurich Information Security and Privacy Center (ZISC). He received his Dipl.Ing. Degree in Electrical Engineering / Computer Science from the University of Split, Croatia, and his Ph.D. degree in Communication Systems from EPFL (Swiss Federal Institute of Technology - Lausanne) in 2004. Prior to joining ETH Zurich in 2006 he was a postdoctoral researcher in the Networked & Embedded Systems Laboratory (NESL), University of California Los Angeles and an Assistant Professor in the Informatics and Mathematical Modeling Department (IMM), Technical University of Denmark (DTU).