

# CRYPTOGRAPHIC CLOUD STORAGE

SENY KAMARA, KRISTIN LAUTER – MICROSOFT RESEARCH CRYPTOGRAPHY GROUP

## 1. Introduction

Advances in networking technology and an increase in the need for computing resources have prompted many organizations to outsource their storage and computing needs. This new economic and computing model is commonly referred to as *cloud computing* and includes various types of services such as: infrastructure as a service (IaaS), where a customer makes use of a service provider's computing, storage or networking infrastructure; platform as a service (PaaS), where a customer leverages the provider's resources to run custom applications; and finally software as a service (SaaS), where customers use software that is run on the provider's infrastructure.

Cloud infrastructures can be roughly categorized as either private or public. In a private cloud, the infrastructure is managed and owned by the customer and located on-premise (i.e., in the customer's region of control). In particular, this means that access to customer data is under its control and is only granted to parties it trusts. In a public cloud the infrastructure is owned and managed by a cloud service provider and is located off-premise (i.e., in the cloud service provider's region of control). This means that customer data is outside its control and could potentially be granted to untrusted parties.

Storage services based on public clouds such as Microsoft's Azure storage service and Amazon's S3 provide customers with scalable and dynamic storage. By moving their data to the cloud customers can avoid the costs of building and maintaining a private storage infrastructure, opting instead to pay a service provider as a function of its needs. For most customers, this provides several benefits including availability (i.e., being able to access data from anywhere) and reliability (i.e., not having to worry about backups) at a relatively low cost.

While the benefits of using a public cloud infrastructure are clear, it introduces significant security and privacy risks. In fact, it seems that the biggest hurdle to the adoption of cloud storage (and cloud computing in general) is concern over the confidentiality and integrity of data. While, so far, consumers have been willing to trade privacy for the convenience of software services (e.g., for web-based email, calendars, pictures etc...), this is not the case for enterprises and government organizations. This reluctance can be attributed to several factors that range from a desire to protect mission-critical data to regulatory obligations to preserve the confidentiality and integrity of data. The latter can occur when the customer is responsible for keeping personally identifiable information (PII), or medical and financial records. So while cloud storage has enormous promise, unless the issues of confidentiality and integrity are addressed many potential customers will be reluctant to make the move.

To address the concerns outlined above and increase the adoption of cloud storage, we argue for designing a *virtual private storage service* based on new cryptographic techniques. Such a service should aim to achieve the "best of both worlds" by providing the security of a private cloud and the functionality and cost savings of a public cloud. More precisely, such a service should provide (at least):

- **confidentiality:** the cloud storage provider does not learn any information about customer data
- **integrity:** any unauthorized modification of customer data by the cloud storage provider can be detected by the customer
- **non-repudiation:** any access to customer data is logged,

while retaining the main benefits of a public storage service:

- **availability:** customer data is accessible from any machine and at all times
- **reliability:** customer data is reliably backed up
- **efficient retrieval:** data retrieval times are comparable to a public cloud storage service
- **data sharing:** customers can share their data with trusted parties.

An important aspect of a cryptographic storage service is that the security properties described above are achieved based on strong cryptographic guarantees as opposed to legal, physical and access control mechanisms. We believe this has several important benefits which we discuss further in Section 3.

This article is organized as follows. In Section 2 we describe, at a high level, a possible architecture for a cryptographic storage service. We consider both consumer and enterprise scenarios. We stress that this design is not intended to be a formal specification (indeed many important business and engineering questions still need to be addressed) but is only meant to serve as an illustration of how some of the new and non-standard cryptographic techniques that have been developed recently can be combined to achieve our goals. In Section 3 we give an overview of the benefits of a cryptographic storage service, e.g., reducing the legal exposure of both customers and cloud providers, and achieving regulatory compliance. In Section 4 we describe in more detail the relevant cryptographic techniques, including searchable encryption, proofs of storage and policy-based encryption. Finally, in Section 5, we mention some cloud services that could be built on top of a cryptographic storage service such as secure back-ups, archival, health record systems, secure data exchange and e-discovery.

## 2. Architecture of a Cryptographic Storage Service

We now describe, at a high level, a possible architecture for a cryptographic storage service. At its core, the architecture consists of three components: a *data processor* (DP), that processes data before it is sent to the cloud; a *data verifier* (DV), that checks whether the data in the cloud has been tampered with; and a *token generator* (TG), that generates tokens which enable the cloud storage provider to retrieve segments of customer data. We describe designs for both consumer and enterprise scenarios.

### A Consumer Architecture

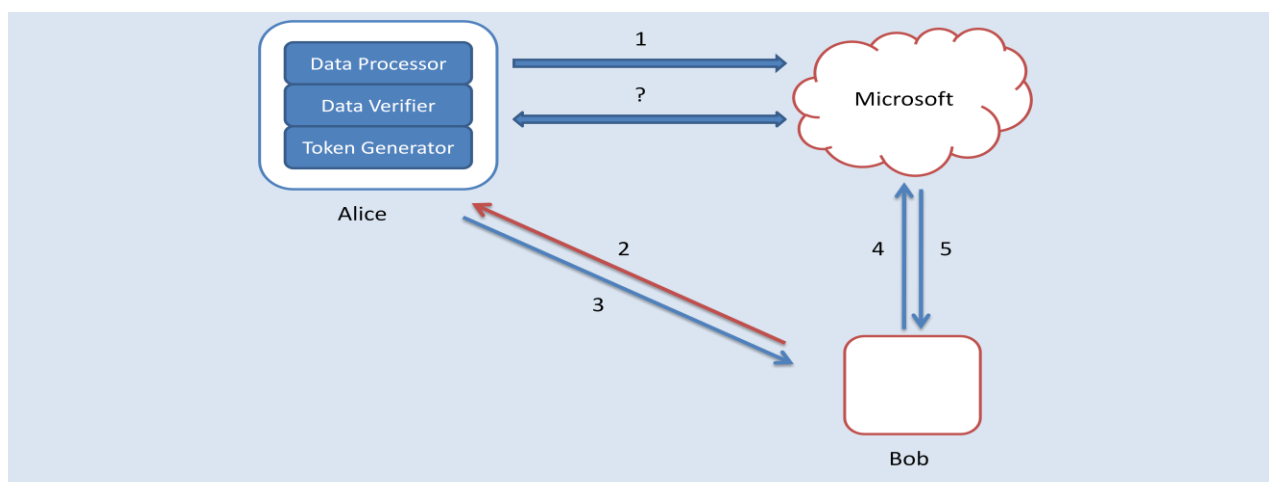
Consider three parties: a user Alice that stores her data in the cloud; a user Bob with whom Alice wants to share data; and a cloud storage provider that stores Alice's data. To use the service, Alice and Bob begin by downloading a client application that consists of a data processor, a data verifier and a token generator. Upon its first execution, Alice's application generates a cryptographic key. We will refer to this key as a master key and assume it is stored locally on Alice's system and that it is kept secret from the cloud storage provider.

Whenever Alice wishes to upload data to the cloud, the data processor is invoked. It attaches some metadata (e.g., current time, size, keywords etc...) and encrypts and encodes the data and metadata with a variety of

cryptographic primitives (which we describe in more detail in Section 4). Whenever Alice wants to verify the integrity of her data, the data verifier is invoked. The latter uses Alice’s master key to interact with the cloud storage provider and ascertain the integrity of the data. When Alice wants to retrieve data (e.g., all files tagged with keyword “urgent”) the token generator is invoked to create a token and a decryption key. The token is sent to the cloud storage provider who uses it to retrieve the appropriate (encrypted) files which it returns to Alice. Alice then uses the decryption key to decrypt the files.

Data sharing between Alice and Bob proceeds in a similar fashion. Whenever she wishes to share data with Bob, the token generator is invoked to create a token and a decryption key which are both sent to Bob. He then sends the token to the provider who uses it to retrieve and return the appropriate encrypted documents. Bob then uses the decryption key to recover the files.

This process is illustrated in Figure 1. We note that in order to achieve the security properties we seek, it is important that the client-side application and, in particular, the core components be either open-source or implemented or verified by someone other than the cloud service provider.



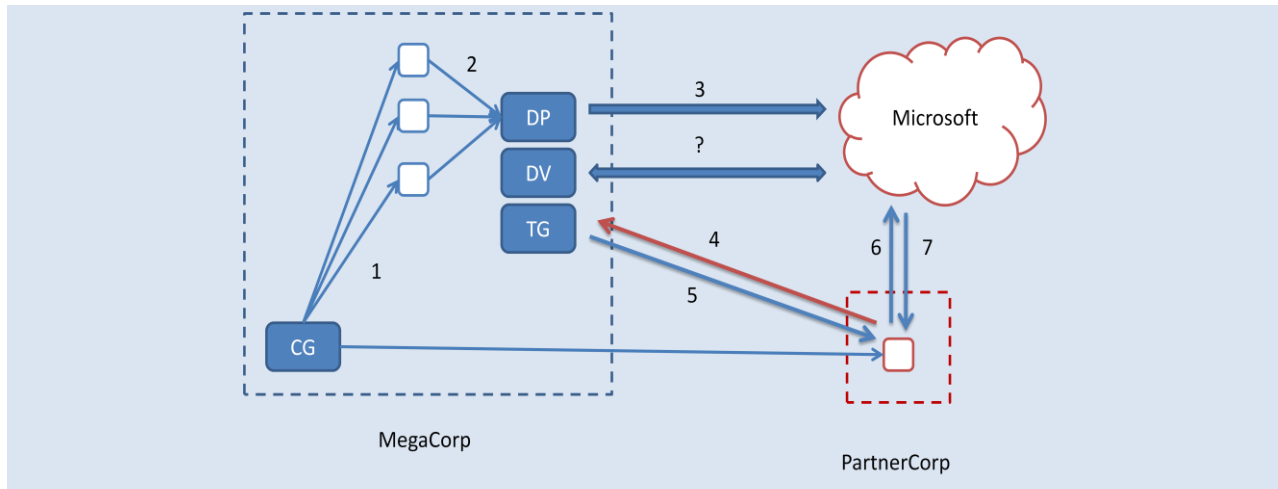
**Figure 1:** (1) Alice’s data processor prepares the data before sending it to the cloud; (2) Bob asks Alice for permission to search for a keyword; (3) Alice’s token generator sends a token for the keyword and a decryption key back to Bob; (4) Bob sends the token to the cloud; (5) the cloud uses the token to find the appropriate encrypted documents and returns them to Bob. At any point in time, Alice’s data verifier can verify the integrity of the data.

## An Enterprise Architecture

In the enterprise scenario we consider an enterprise MegaCorp that stores its data in the cloud; a business partner PartnerCorp with whom MegaCorp wants to share data; and a cloud storage provider that stores MegaCorp’s data. To handle enterprise customers, we introduce an extra component: a *credential generator*. The credential generator implements an access control policy by issuing credentials to parties inside and outside MegaCorp.

To use the service, MegaCorp deploys dedicated machines within its network. Depending on the particular scenario, these dedicated machines will run various core components. Since these components make use of a master secret key, it is important that they be adequately protected and, in particular, that the master key be kept secret from the cloud storage provider and PartnerCorp. If this is too costly in terms of resources or expertise, management of the dedicated machines (or specific components) can alternatively be outsourced to a trusted entity.

In the case of a medium-sized enterprise with enough resources and expertise, the dedicated machines include a data processor, a data verifier, a token generator and a credential generator. To begin, each MegaCorp and PartnerCorp employee receives a credential from the credential generator. These credentials will reflect some relevant information about the employees such as, e.g., their organization or team or role. Whenever a MegaCorp employee



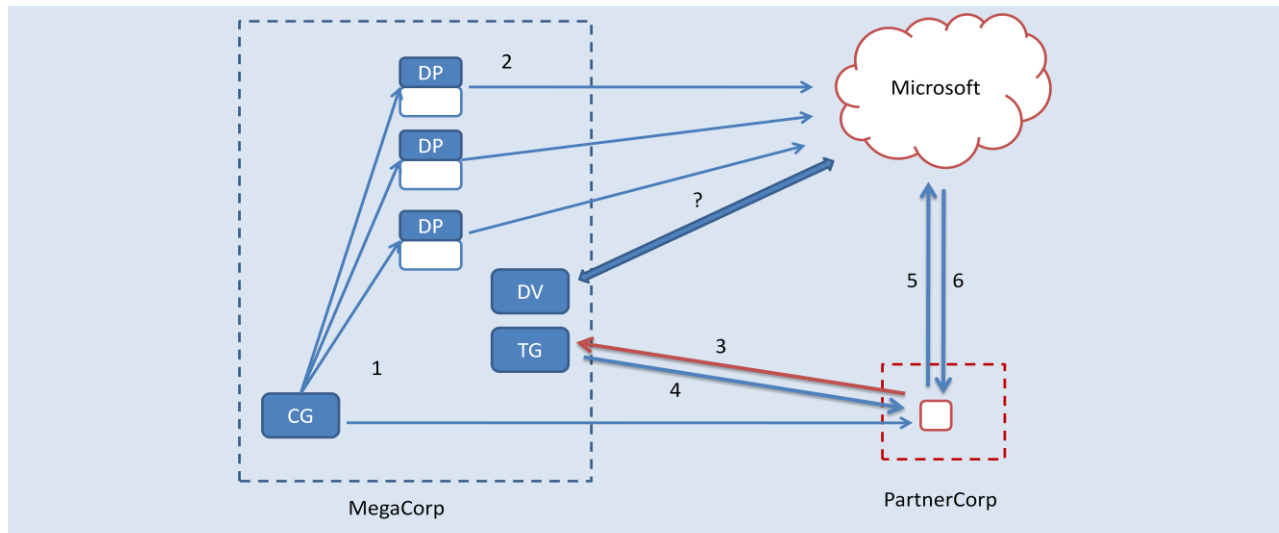
**Figure 2:** (1) Each MegaCorp and PartnerCorp employee receives a credential; (2) MegaCorp employees send their data to the dedicated machine; (3) the latter processes the data using the data processor before sending it to the cloud; (4) the PartnerCorp employee sends a keyword to MegaCorp's dedicated machine ; (5) the dedicated machine returns a token; (6) the PartnerCorp employee sends the token to the cloud; (7) the cloud uses the token to find the appropriate encrypted documents and returns them to the employee. At any point in time, MegaCorp's data verifier can verify the integrity of MegaCorp's data.

generates data that needs to be stored in the cloud, it sends the data together with an associated decryption policy to the dedicated machine for processing. The decryption policy specifies the type of credentials necessary to decrypt the data (e.g., only members of a particular team). To retrieve data from the cloud (e.g., all files generated by a particular employee), an employee requests an appropriate token from the dedicated machine. The employee then sends the token to the cloud provider who uses it to find and return the appropriate encrypted files which the employee decrypts using his credentials. Whenever MegaCorp wants to verify the integrity of its data, the dedicated machine's data verifier is invoked. The latter uses the master secret key to interact with the storage provider and ascertain the integrity of the data.

Now consider the case where a PartnerCorp employee needs access to MegaCorp's data. The employee authenticates itself to MegaCorp's dedicated machine and sends it a keyword. The latter verifies that the particular search is allowed for this PartnerCorp employee. If so, the dedicated machine returns an appropriate token which the employee uses to recover the appropriate files from the service provider. It then uses its credentials to decrypt the file. This process is illustrated in Figure 2. Similarly to the consumer architecture, it is imperative that all components be either open-source or implemented by someone other than the cloud service provider.

In the case that MegaCorp is a very large organization and that the prospect of running and maintaining enough dedicated machines to process all employee data is infeasible, consider the following slight variation of the architecture described above. More precisely, in this case the dedicated machines only run data verifiers, token generators and credential generators while the data processing is distributed to each employee. This is illustrated in Figure 3. Note that in this scenario the data processors do not include the master secret key so the confidentiality of the data is not affected. The data processors, however, do include some keying material which, if

revealed to the service provider, could enable it to compromise the confidentiality of the tokens it receives (i.e., it could learn which keywords are being searched for).



**Figure 3:** (1) Each MegaCorp and PartnerCorp employee receives a credential; (2) MegaCorp employees process their data using their own data processors and send them to the cloud; (3) the PartnerCorp employee sends a keyword to MegaCorp’s dedicated machine; (4) the latter returns a token; (5) the employee sends the token to the cloud; (6) the cloud uses the token to find the appropriate encrypted documents and returns them to the employee. At any point in time, MegaCorp’s data verifier can the integrity of MegaCorp’s data.

### 3. Benefits of a Cryptographic Storage Service

The core properties of a cryptographic storage service are that (1) control of the data is maintained by the customer and (2) the security properties are derived from cryptography, as opposed to legal mechanisms, physical security or access control. Therefore, such a service provides several compelling advantages over other storage services based on public cloud infrastructures. In this section, we recall some of the main concerns with cloud computing as outlined in the Cloud Security Alliance’s recent report (Cloud Security Alliance 2009) and highlight how these concerns can be mitigated by such an architecture.

**Regulatory compliance.** Most countries have laws in place that make organizations responsible for the protection of the data that is entrusted to them. This is particularly so for the case of personally identifiable information, medical records and financial records. And since organizations are often held responsible for the actions of their contractors the use of a public cloud storage service can involve significant legal risks. In a cryptographic storage service, the data is encrypted on-premise by the data processor(s). This way, customers can be assured that the confidentiality of their data is preserved irrespective of the actions of the cloud storage provider. This greatly reduces any legal exposure for both the customer and the provider.

**Geographic restrictions.** Data that is stored in certain legal jurisdictions may be subject to regulations even if it was not collected there. Because it can be difficult to ascertain exactly where one’s data is being stored once it is sent to the cloud (i.e., many service providers have data centers deployed throughout the world) some customers may be reluctant to use a public cloud for fear of increasing their legal exposure. In a cryptographic storage service data

is only stored in encrypted form so any law that pertains stored data has little to no effect on the customer. This reduces legal exposure for the customer and allows the cloud storage provider to make optimal use of its storage infrastructure, thereby reducing costs.

**Subpoenas.** If an organization becomes the subject of an investigation, law enforcement agencies may request access to its data. If the data is stored in a public cloud, the request may be made to the cloud provider and the latter could even be prevented from notifying the customer. This can have severe consequences for customers. First, it preempts the customer from challenging the request. Second, it can lead to law enforcement having access to data from clients that are not under investigation (Wired 2009). Such a scenario can occur due to the fact that service providers often store multiple customer's data on the same disks. In a cryptographic storage service, since data is stored in encrypted form and since the customer retains possession of all the keys, any request for the data must be made directly to the customer.

**Security breaches.** Even if a cloud storage provider implements strong security practices there is always the possibility of a security breach. If this occurs the customer may be legally responsible. In a cryptographic storage service data is encrypted and data integrity can be verified at any time. Therefore, a security breach poses little to no risk for the customer.

**Electronic discovery.** Digital information plays an important role in legal proceedings and often organizations are required to preserve and produce records for litigation. Organizations with high levels of litigation may need to keep a copy of large amounts of data on-premise in order to assure its integrity. This can obviously negate the benefits of using a cloud storage service. Since, with a cryptographic storage service, a customer can verify the integrity of its data at any point in time (e.g., every hour) a provider has every incentive to preserve its integrity.

**Data retention and destruction.** In many cases a customer may be responsible for the retention and destruction of data it has collected. If this data is stored in the cloud, however, it can be difficult for a customer to ascertain the integrity of the data or to verify whether it was properly discarded. A cryptographic storage service alleviates these concerns since data integrity can be verified and since the information necessary to decrypt data (i.e., the master key) is kept on-premise. Secure data erasure can be effectively achieved by just erasing the master key.

## 4. Implementing the Core Components

The core components of a cryptographic storage service can be implemented using a variety of techniques, some of which were developed specifically for cloud computing. When preparing data for storage in the cloud, the data processor begins by indexing it and encrypting it with a symmetric encryption scheme (e.g., AES) under a unique key. It then encrypts the index using a *searchable encryption* scheme and encrypts the unique key with an *attribute-based encryption* scheme under an appropriate policy. Finally, it encodes the encrypted data and index in such a way that the data verifier can later verify their integrity using a *proof of storage*.

In the following we provide high level descriptions of these new cryptographic primitives. While traditional techniques like encryption and digital signatures could be used to implement the core components, they would do so at considerable cost in communication and computation. To see why, consider the example of an organization that encrypts and signs its data before storing it in the cloud. While this clearly preserves confidentiality and integrity it has the following limitations. To enable searching over the data, the customer has to either store an index locally, or download all the (encrypted) data, decrypt it and search locally. The first approach obviously negates the benefits of cloud storage (since indexes can grow large) while the second scales poorly. With respect

to integrity, note that the organization would have to retrieve all the data first in order to verify the signatures. If the data is large, this verification procedure is obviously undesirable. Various solutions based on (keyed) hash functions could also be used, but all such approaches only allow a fixed number of verifications.

## Searchable Encryption

A searchable encryption scheme provides a way to encrypt a search index so that its contents are hidden except to a party that is given appropriate tokens. More precisely, consider a search index generated over a collection of files (this could be a full-text index or just a keyword index). Using a searchable encryption scheme, the index is encrypted in such a way that (1) given a token for a keyword one can retrieve pointers to the encrypted files that contain the keyword; and (2) without a token the contents of the index are hidden. In addition, the tokens can only be generated with knowledge of a secret key and the retrieval procedure reveals nothing about the files or the keywords except that the files contain a keyword in common.

This last point is worth discussing further as it is crucial to understanding the security guarantee provided by searchable encryption. Notice that over time (i.e., after many searches) knowing that a certain subset of documents contain a word in common *may* leak some useful information. This is because the server could make some assumptions about the client's search pattern and use this information to make a guess about the keywords being searched for. It is important to understand, however, that while searching does leak *some* information to the provider, what is being leaked is exactly what the provider would learn from the act of returning the appropriate files to the customer (i.e., that these files contain *some* keyword in common). In other words, the information "leaked" to the cloud provider is not leaked by the cryptographic primitives, but by the manner in which the service is being used (i.e., to fetch files based on exact keyword matches). It is important to understand that this leakage is in some sense inherent to any efficient and reliable cloud storage service and is, at worst, less information than what is leaked by using a public cloud storage service. The only known alternative, which involves making the service provider return false positives and having the client perform some local filtering, is inefficient in terms of communication and computational complexity.

There are many types of searchable encryption schemes, each one appropriate to particular application scenarios. For example, the data processors in our consumer and small enterprise architectures use symmetric searchable encryption (SSE), while the data processors in our large enterprise architecture uses asymmetric searchable encryption (ASE). In the following we describe each type of scheme in more detail.

**Symmetric searchable encryption.** SSE is appropriate in any setting where the party that searches over the data is also the one who generates it. Borrowing from storage systems terminology, we refer to such scenarios as single writer/single reader (SWSR). SSE schemes were introduced in (Song, Wagner and Perrig 2000) and improved constructions and security definitions were given in (Goh 2003), (Chang and Mitzenmacher 2005) and (Curtmola, et al. 2006).

The main advantages of SSE are efficiency and security while the main disadvantage is functionality. SSE schemes are efficient both for the party doing the encryption and (in some cases) for the party performing the search. Encryption is efficient because most SSE schemes are based on symmetric primitives like block ciphers and pseudo-random functions. Search can be efficient because the typical usage scenarios for SSE (i.e., SWSR) allow the data to be pre-processed and stored in efficient data structures.

The security guarantees provided by SSE are, roughly speaking, the following:

- (1) without any trapdoors the server learns nothing about the data except its length

(2) given a trapdoor for a keyword  $W$ , the server learns which (encrypted) documents contain  $W$  without learning  $W$ .

While these security guarantees are stronger than the ones provided by both asymmetric and efficiently searchable encryption (described below), we stress that they do have their limitations (as described above).

The main disadvantage of SSE is that the known solutions tradeoff efficiency and functionality. This is easiest to see by looking at two of the main constructions proposed in the literature. In the scheme proposed by Curtmola et al. (Curtmola, et al. 2006), search time for the server is optimal (i.e., linear in the number of documents that contain the keyword) but updates to the index are inefficient. On the other hand, in the scheme proposed by Goh (Goh 2003), updates to the index can be done efficiently but search time for the server is slow (i.e., linear in the total number of documents). We also remark that neither scheme handles searches that are composed of conjunctions or disjunction of terms. The only SSE scheme that handles conjunctions (Golle, Waters and Staddon 2004) is based on pairings on elliptic curves and is as inefficient as the asymmetric searchable encryption schemes discussed below.

**Asymmetric searchable encryption (ASE)<sup>1</sup>**. ASE schemes are appropriate in any setting where the party searching over the data is different from the party that generates it. We refer to such scenarios as many writer/single reader (MWSR). ASE schemes were introduced in (Boneh, Di Crescenzo, et al. 2004). Improved definitions were proposed in (Abdalla, et al. 2005) and schemes that handle conjunctions were given in (Park, Kim and Lee 2005) and (Boneh and Waters, Conjunctive, Subset, and Range Queries on Encrypted Data 2007).

The main advantage of ASE is functionality while the main disadvantages are inefficiency and weaker security. Since the writer and reader can be different, ASE schemes are usable in a larger number of settings than SSE schemes. The inefficiency comes from the fact that all known ASE schemes require the evaluation of pairings on elliptic curves which is a relatively slow operation compared to evaluations of (cryptographic) hash functions or block ciphers. In addition, in the typical usage scenarios for ASE (i.e., MWSR) the data cannot be stored in efficient data structures.

The security guarantees provided by ASE are, roughly speaking, the following:

- (1) without any trapdoors the server learns nothing about the data except its length
- (2) given a trapdoor for a keyword  $W$ , the server learns which (encrypted) documents contain  $W$ .

Notice that (2) here is weaker than in the SSE setting. In fact, when using an ASE scheme, the server can mount a dictionary attack against the token and figure out which keyword the client is searching for (Byun, et al. 2006). It can then use the token (for which it now knows the underlying keyword) and do a search to figure out which documents contain the (known) keyword. Note that this should not necessarily be interpreted as saying that ASE schemes are insecure, just that one has to be very careful about the particular usage scenario and the types of keywords and data being considered.

**Efficient ASE (ESE)**. ESE schemes are appropriate in any setting where the party that searches over the data is different from the party that generates it *and* where the keywords are hard to guess. This falls into the MWSR scenario as well. ESE schemes were introduced in (Bellare, Boldyreva and O' Neill 2007).

---

<sup>1</sup> ASE schemes are sometimes referred to as “public-key encryption schemes with keyword search” (PEKS).

The main advantage of efficient ASE is that search is more efficient than (plain) ASE. The main disadvantage, however, is that ESE schemes are also vulnerable to dictionary attacks. In particular, the dictionary attacks against ESE can be performed directly against the encrypted index (as opposed to against the *token* like in ASE).

**Multi-user SSE (mSSE).** mSSE schemes are appropriate in any setting where many parties wish to search over data that is generated by a single party. We refer to such scenarios as single writer/many reader (SWMR). Multi-user SSE schemes were introduced in (Curtmola, et al. 2006).

In an mSSE scheme, in addition to being able to encrypt indexes and generate tokens, the owner of the data can also add and revoke users' search privileges over his data.

## Attribute-based Encryption

Another set of cryptographic techniques that has emerged recently allows the specification of a decryption policy to be associated with a ciphertext. More precisely, in a (ciphertext-policy) attribute-based encryption scheme each user in the system is provided with a decryption key that has a set of attributes associated with it (this is how the "credentials" in Section 2 would be implemented). A user can then encrypt a message under a public key and a policy. Decryption will only work if the attributes associated with the decryption key match the policy used to encrypt the message. Attributes are qualities of a party that can be established through relevant credentials such as being a Microsoft employee or living in Washington State.

Attribute-based encryption was introduced in (Sahai and Waters 2005). Improved constructions are given in (Goyal, et al. 2006), (Ostrovsky, Sahai and Waters 2007), (Bethencourt, Sahai and Waters 2007) and (Waters 2008). The setting where attributes can be distributed by multiple parties is considered in (Chase, Multi-authority Attribute Based Encryption 2007) and (Chase and Chow, Improving Privacy and Security in Multi-Authority Attribute-Based Encryption 2009).

## Proofs of Storage

A proof of storage is a protocol executed between a client and a server with which the server can prove to the client that it did not tamper with its data. The client begins by encoding the data before storing it in the cloud. From that point on, whenever it wants to verify the integrity of the data it runs a proof of storage protocol with the server. The main benefits of a proof of storage are that (1) they can be executed an arbitrary number of times; and (2) the amount of information exchanged between the client and the server is extremely small and independent of the size of the data.

Proofs of storage can be either privately or publicly verifiable. Privately verifiable proofs of storage only allow the client (i.e., the party that encoded the file) to verify the integrity of the data. With a publicly verifiable proof of storage, on the other hand, anyone that possesses the client's public key can verify the data's integrity.

Proofs of storage were introduced in (Ateniese, Burns, et al. 2007) and (Juels and Kaliski 2007). Improved definitions and constructions were given in (Shacham and Waters 2008), (Ateniese, Kamara and Katz, Proofs of Storage from Homomorphic Identification Protocols 2009) and (Erway, et al. 2009).

## 5. Cloud Services

### Secure Extranet

In addition to simple storage, many enterprise customers will have a need for some associated services. These services can include any number of business processes including sharing of data among trusted partners, litigation support, monitoring and compliance, back-up, archive and audit logs. We refer to a cryptographic storage service together with an appropriate set of enterprise services as a secure extranet and believe this could provide a valuable service to enterprise customers.

### Electronic Health Records

In February 2009, 19 billion dollars were provisioned by the U.S. government to digitize health records. This move towards electronic health records promises to reduce medical errors, save lives and decrease the cost of healthcare. Given the importance and sensitivity of health-related data, it is clear that any storage platform for health records will need to provide strong confidentiality and integrity guarantees to patients and care givers (see (Benaloh, et al. 2009) for more regarding these issues).

### Interactive Scientific Publishing

As scientists continue to produce large data sets which have broad value for the scientific community, demand will increase for a storage infrastructure to make such data accessible and sharable. To incent scientists to share their data, scientific societies such as the Optical Society of America are considering establishing a publication forum for data sets in partnership with industry. Such an interactive publication forum will need to provide strong guarantees to authors on how their data sets may be accessed and used by others, and could be built on a cryptographic cloud storage system like the one proposed here.

## Bibliography

Abdalla, Michel, et al. «Searchable Encryption Revisited: Consistency Properties, Relation to Anonymous IBE, and Extensions.» *Advances in Cryptology -- CRYPTO '05*. Springer, 2005. 205-222.

Ateniese, Giuseppe, et al. «Provable data possession at untrusted stores.» *ACM Conference on Computer and Communications Security*. ACM Press, 2007. 598-609.

Ateniese, Giuseppe, Seny Kamara, et Jonathan Katz. «Proofs of Storage from Homomorphic Identification Protocols.» *Advances in Cryptology -- Asiacrypt '09*. Springer, 2009.

Bellare, Mihir, Alexandra Boldyreva, et Adam O' Neill. «Deterministic and Efficiently Searchable Encryption.» *Advances in Cryptology - CRYPTO 2007*. Springer, 2007. 535-552.

Benaloh, Josh, Melissa Chase, Eric Horvitz, et Kristin Lauter. «Patient Controlled Encryption: Ensuring Privacy of Electronic Medical Records.» *The ACM Cloud Computing Security Workshop*. 2009.

Boneh, Dan, et Brent Waters. «Conjunctive, Subset, and Range Queries on Encrypted Data.» *Theory of Cryptography Conference*. Springer, 2007. 535-554.

Boneh, Dan, Giovanni Di Crescenzo, Rafail Ostrovsky, et Giuseppe Persiano. «Public Key Encryption with Keyword Search .» *Advances in Cryptology - EUROCRYPT 2004*. Springer, 2004. 506-522.

Byun, Jin Wook, Hyun Suk Rhee, Hyun-A Park, et Dong Hoon Lee. «Off-Line Keyword Guessing Attacks on Recent Keyword Search Schemes over Encrypted Data.» *Secure Data Management* . Springer, 2006. 75-83.

Chang, Yan-Cheng, et Michael Mitzenmacher. «Privacy Preserving Keyword Searches on Remote Encrypted Data.» *Applied Cryptography and Network Security*. Springer, 2005. 442-455.

Chase, Melissa. «Multi-authority Attribute Based Encryption.» *Theory of Cryptography Conference*. Springer, 2007. 515-534.

Chase, Melissa, et Sherman Chow. «Improving Privacy and Security in Multi-Authority Attribute-Based Encryption.» *ACM Conference on Computer and Communications Security*. 2009.

Cloud Security Alliance. «Security Guidance for Critical Areas of Focus in Cloud Computing.» April 2009.  
<http://www.cloudsecurityalliance.org/guidance/csaguide.pdf>.

Curtmola, Reza, Juan Garay, Seny Kamara, et Rafail Ostrovsky. «Searchable symmetric encryption: improved definitions and efficient constructions.» *13th ACM Conference on Computer and Communications Security (CCS)*. ACM Press., 2006. 79-88.

Erway, Chris, Alptekin Kupcu, Charalampos Papamanthou, et Roberto Tamassia. «Dynamic Provable Data Possession.» *ACM Conference on Computer and Communications Security*. 2009.

Goh, E.-J. *Secure Indexes*. IACR ePrint, 2003.

Golle, Phillipe, Brent Waters, et Jessica Staddon. «Secure Conjunctive Keyword Search over Encrypted Data.» *Applied Cryptography and Network Security (ACNS '04)*. 2004. 31-45.

Goyal, Vipul, Omkant Pandey, Amit Sahai, et Brent Waters. «Attribute-Based Encryption for Fine-Grained Access Control of Encrypted Data.» *ACM Conference on Computer and Communications Security*. 2006. 89-98.

Juels, Ari, et Burt Kaliski. «PORs: proofs of retrievability for large files.» *ACM Conference on Computer and Communications Security*. ACM Press, 2007. 584-597.

Ostrovsky, Rafail, Amit Sahai, et Brent Waters. «Attribute-based encryption with non-monotonic access structures.» *ACM Conference on Computer and Communications Security*. 2007. 195-203.

Park, Dong Jin, Kihyun Kim, et Pil Joong Lee. «Public Key Encryption with Conjunctive Field Keyword Search.» *Information Security Applications*. Springer, 2005. 73-86.

Sahai, Amit, et Brent Waters. «Fuzzy Identity-Based Encryption. .» *Advances in Cryptology -- Eurocrypt '05*. Springer, 2005. 457-473.

Shacham, Hovav, et Brent Waters. «Compact Proofs of Retrievability.» *ASIACRYPT*. Springer, 2008. 90-107.

Song, Dawn, David Wagner, et Adrian Perrig. «Practical Techniques for Searches on Encrypted Data.» *IEEE Symposium on Security and Privacy*. IEEE Press, 2000. 44-55.

Wired. «Company Caught in Texas Data Center Raid Loses Suit Against FBI.» 8 April 2009.  
<http://www.wired.com/threatlevel/2009/04/company-caught/>.