

Best Practices for Cloud Data Protection and Key Management



Introduction

Digital transformation is a persistent trend that has resulted in fundamental shift in storage, access, and management of digital assets. This shift is particularly noticeable for the most prized asset: an organization's sensitive data. Sensitive data used to reside in controlled on premises enterprise data centers protected by both physical and logical security controls. Although, until recently it was unthinkable to trust a third-party with such data, organizations are now accelerating the migration of their workloads and data to the cloud. According to the Thales 2019 Global Data Threat Report¹, over 40% of organizations now store their sensitive data using infrastructure (IaaS), platform (PaaS) or software (SaaS) services of public Cloud Service Providers (CSPs), such as Amazon Web Services, Microsoft Azure, and Google Cloud. This change, though initially driven by economics, has delivered greater flexibility and elasticity in management of computing and storage resources. A clear casualty of this change has been the traditional method of data security that relied on a well-defined network perimeter to protect data.

Once particularly attractive to startups and small to medium size enterprises, cloud has now become a global transformation embraced by all enterprises. At the same time consumers expect data and services to be available from anywhere, at any time, and from any device. Anything less degrades user experience creating a drag on productivity. No wonder businesses look to the cloud to meet such demanding requirements. However, amid this rush to the cloud, data security is often neglected². Cloud deployments can considerably increase the attack surface, especially in case of multi-cloud use, providing more options for fraudsters to get access to sensitive assets³. Data in the cloud cannot be protected by protecting a single known network perimeter; the cloud introduces new cybersecurity challenges, because the environments themselves in which sensitive data resides may no longer be trusted.

One example of sensitive data is user personal identifiable information (PII). Protecting PII and consumer privacy is making headline news these days. Politicians, government officials and privacy advocates are demanding that all entities holding customer data be held accountable for protecting the data and consequently the privacy of end customers. Data privacy regulations, such as the European Union's General Data Protection Regulation (GDPR)⁴ and the California Consumer Protection Act (CCPA)⁵, require strict compliance and impose stiff penalties for data breaches that result in loss of consumer sensitive data and privacy. Digital transformation must be done within the context of such regulations.

Recognizing the importance of data protection, cloud service providers now offer data encryption and key management services. These services can be used across the different types of infrastructure offered by the respective service providers. However, while the native encryption and key management services offer good-enough protection, many organizations, especially those in highly regulated industries, such as finance, banking, insurance and health care, need higher levels of assurance for risk management and compliance. An important tool for risk management is the ability for an organization to be the custodian of data encryption keys instead of delegating key ownership to service providers. Many security conscious organizations have an additional requirement that all data encryption keys must be created, stored, and managed using a FIPS 140-2 level 3 certified key manager. In response to such higher assurance requirements, cloud service providers have begun offering features, such as Bring Your Own Key (BYOK) and Hold Your Own Key (HYOK).

This paper describes security best practices for protecting sensitive data in the public cloud, and explains concepts such as BYOK, HYOK, key brokering, and Root of Trust (RoT). It explains the level of data protection that can be achieved by using cloud native encryption and key management services, and how these can be augmented by allowing customers to take more responsibility for, and control over, their keys.

1 Thales 2019 Global Data Threat Report, <https://cpl.thalesgroup.com/resources/encryption/2020/data-threat-report>

2 2019 Thales Cloud Security Report, <https://cpl.thalesgroup.com/resources/cloud-security/2019/thales-global-cloud-security-study-white-paper>

3 Gruschka, Nils & Jensen, Meiko. (2010). Attack Surfaces: A Taxonomy for Attacks on Cloud Services. 276 - 279. 10.1109/CLOUD.2010.23.

4 Rules for the protection of personal data inside and outside of EU, European Commission, https://ec.europa.eu/info/law/law-topic/data-protection_en

5 California Consumer Privacy Act (CCPA), <https://oag.ca.gov/privacy/ccpa>

Security Principles

It is critical to adopt business practices that strike the right balance between security of data, assets, and user identities, and the cost associated with offering such security.

Security Basics

Since data is no longer confined within traditional network perimeters, it can reside in a diverse set of environments during its lifecycle from creation to deletion. Each environment brings its own challenges that have to be addressed using specific technologies. Three of the more commonly considered environments are:

- **At-Rest.** Data is predominantly in this “inactive” state, when it is stored physically in any digital form (e.g., databases, files, tape-backups, etc.) and remains unchanged.
- **In-Use.** In this state, data is “active” and being manipulated (e.g., operation on a row in a database, processing of data in a CPU, etc.). Most often, data protection mechanisms applied during at-rest state are removed before data is operated upon. However, there are emerging technologies that can also prevent exposure of data during its use⁶.
- **In-Transit.** This state implies that data is transported from one physical location or medium to another; e.g., transmitted across the network or transferred from persistent memory to the CPU, etc.

Data security principles require that data is protected in each of these states. Security controls must defend against all threats applicable to the medium, including side-channel attacks, where information is gleaned from valid operation of the device, rather than from exploitation of a weakness. Table 1 summarizes the building blocks of data protection.

Table 1: Data protection fundamentals

Concept	Description
Confidentiality	Keep information secret and private
Integrity	Prevent unauthorized changes to information
Availability	Keep system/data available for use
Accountability	Hold users/system accountable for actions
Auditability	Keep verifiable records of activities

Separation of Duties

It is essential to ensure that fundamental principles of both separation of duties and least privileged access are followed for data protection. This involves enforcing separation of duties between entities processing and storing data from those providing security services. In other words, ideally, the data storage provider and the encryption service provider should be separate entities; a distinction that is often overlooked. For example, cloud storage providers do encrypt data both during transmission and before storage. However, since they also hold the encryption key for stored data, they have direct access to all data that resides on their servers. For enterprises that own the data, this provides little comfort since trust has to be placed entirely in the hands of cloud storage providers. Such encryption merely protects against data breaches in case attackers get access to the encrypted data stored on the server. However, even in this case, there is a possibility that if data is breached due to a configuration error or a security flaw at the cloud storage provider, the keys that encrypt the data may also be compromised.

As such, a growing number of enterprises have started to look into solutions that encrypt data and files before sending them to cloud storage providers. Similarly, cloud service providers advocate security, in particular data security, as being a shared responsibility between the customer and the storage provider. Figure 1 illustrates this shared responsibility where CSP handles security of infrastructure and the customer handles security of data.

⁶ Victor Costan and Srinivas Devadas. “Intel SGX Explained.” IACR Cryptology ePrint Archive 2016 (2016): 86. <https://eprint.iacr.org/2016/086.pdf>



Separation of duty for data stored in the cloud

Figure 1: Shared responsibility model for cloud storage

The customer responsibility can be achieved in two ways:

- **Third-Party encryption.** In this model a third-party server encrypts data before sending it to the cloud storage provider. This encryption server sits between the customer and the storage provider, like a Cloud Access Security Broker (CASB). In case encrypted data is breached at the storage provider, no damage is done, since keys are kept in a separate server and cannot be accessed.
- **Data-Origin encryption.** In this model encryption is done entirely on the customer side as close as possible to the origin of data, for example by using an agent in the browser, deploying file/folder encryption or in an application that generates the data. Several cloud storage providers support this model^{7 8}.

Authentication

Authentication is a mechanism of determining if a person requiring access to a resource is the same person he or she claims to be. This is a basic requirement for any system that manages access control to application or data resources. Methods to verify the authenticity of a user can be divided into three broad categories: what-you-know, what-you-have, and what-you-are. Any combination of two or more of these categories is referred to as multi-factor authentication (MFA). Good data protection demands that access to both data and keys be controlled by stronger authentication methods.

The what-you-know category represents knowledge-based authentication and is the oldest, and most widely used, but also the weakest method of identifying users. An example of this is the password. The what-you-have techniques base their security on possession of a unique hardware token, such as a smart card, Universal Serial Bus (USB) token, or a mobile device. This device is paired with the user and can then perform some cryptographic computation on behalf of the user to attest his identity. Examples of such devices are One Time Password (OTP) tokens or Public Private Infrastructure (PKI) tokens, such as smart cards. Finally, what-you-are, places its security on the natural habits or biological characteristics of the user. The user does not have to remember any secret nor carry any dedicated device. Examples of this category can be biometric characteristics, such as a fingerprint, iris scan, or facial recognition.

Access Management

Access Management (AM) is the next phase of creating a robust security framework that builds on strong authentication. It adds fine-grained control to all user interactions with a system. For example, instead of authenticating credentials only once at the perimeter, a user's security posture should be continuously monitored using checks, such as:

- What applications and data does the user have access to? This sets the high-level authorization scope for the user, which is further refined dynamically based on additional AM policy rules.
- For how long is the access granted?
- From what locations is the access allowed?
- During what times is the access allowed?

The goal is to manage access to data and restricted resources based on policies that are set by an administrator and enforced by real-time or near real-time monitoring of user activities.

7 Amazon Simple Storage Service (S3), <https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingClientSideEncryption.html>

8 Azure data encryption at rest, <https://docs.microsoft.com/en-us/azure/security/fundamentals/encryption-atrest>

Data Encryption

Many CSPs offer encryption and key management that are used with their own compute, storage, database, and other services. Choosing the right type of encryption and key management with each of the services requires a deep understanding of the service providers' service offerings. Furthermore, one must understand the level of risk associated with the choice of an encryption and key management approach. Each of the service providers implements encryption and key management differently, but many of them offer customers a choice of being the custodian of keys. Most providers refer to this service using some form of the expression Key Management Service (KMS). KMS enables not only creation and management of keys but also provides control over the use of these keys in services and applications supported by the CSP.

Table 2 summarizes the encryption and key management services of major cloud providers. It is worth noting that there are two types of encryption models offered:

- **Server-side encryption:** The encryption is done by the service provider using encryption keys that are either:
 - Created and managed by the CSP, transparent to the user
 - Customer provided keys that are managed by the CSP
 - Customer hosted keys stored in customer's repositories on their premises and outside the control of CSP
- **Client-side encryption:** Data is encrypted in the customer data center, or by a customer-provisioned application running in the service provider infrastructure, using customer-provisioned keys. In this model, the service providers do not have access to the encryption keys and therefore cannot decrypt the data..

Table 2: Some data protection offerings from cloud providers with service examples

CSP	Service	Data Encryption Option	Key Mgmt.	Comments
Amazon Web Services	S3	Server side, Client side	AWS KMS	Encryption keys can be provided by customer or by KMS
	EC2	Server side, Client side		
	RDS	Server side, Client side		
Microsoft Azure	Virtual Machines, Containers	Disk encryption	Key Vault	Transparent data encryption for server side. Keys can be provided by customer.
	Databases: SQL Redis	Server side, Client side		
Google Cloud	Storage, Compute Engine, Platform	Transparent multi-layered encryption	Cloud KMS	Allows customer provided and customer hosted keys

Root of Trust

Root of Trust (RoT) refers to a source that can be trusted within a cryptographic system. Because cryptographic security is dependent on the actual security of the keys used to encrypt and decrypt data and perform functions, such as generating digital signatures and verifying signatures, RoT schemes generally include hardened and trusted physical devices.

Entropy

Modern cryptography techniques rely on keys to mathematically transform data. These techniques require that keys are unique, have a low probability of collision, and are not predictable. Consequently, the key generation algorithms require a high degree of randomness, known as high entropy. Since keys transform data so that confidentiality is maintained even in the event of data exposure, they are of exceptionally high value. Keys thereby require special handling to ensure only authorized people and systems use them. This service is provided through specialized and hardened devices.

Hardware Security Module

A Hardware Security Module (HSM) is a hardened physical computing device that specializes in key creation, key storage, and execution of cryptographic processes, thereby acting as RoT and bringing the much needed entropy to cryptographic operations. These modules can be FIPS 140-2⁹ certified and traditionally come in the form of a plug-in card that is inserted directly into a device, or as a network accessible appliance. In both cases, the goal is to deliver a trusted computing environment that protects confidential materials such as keys and uses these keys to securely perform cryptographic operations on sensitive data.

The need for such dedicated devices arises, because typical servers cannot guarantee conformance to the aforementioned goals for trust, entropy, and performance. For example, an application running on a server, which may also be providing keys, may contain a database access code, user interface code, business-logic, and cryptographic calls. For mission critical applications this is a dangerous approach, as an attacker can use exposed application endpoints to access cryptographic material, steal keys, and even alter output of cryptographic operations. To prevent such a scenario applications must be segregated into two parts. One part does business logic, database access, and user interface. The second part does the more sensitive tasks such as key storage and cryptographic operation. It is this second part that should run inside the HSM. Trust can be achieved, because an HSM typically offers most of the following security features:

- It is built using specialized hardware that is manufactured in controlled facilities, is well tested, and is certified.
- It has a security-focused OS that ensures no virus or malware is present.
- It has limited and controlled access (both physical and logical) via network interfaces that can be managed by internal firewall and network rules.
- It actively hides and protects cryptographic material, so that this sensitive material cannot be accessed by probing the silicon.
- It offers true, high entropy random number generation, which forms the foundation of secure cryptographic operation.

In addition to incorporating these security safeguards, HSMs are built for performance. Since they are dedicated appliances, they can perform cryptographic operations faster than more general purpose servers. Furthermore, several providers of HSMs have also started to evaluate support for post-quantum cryptography¹⁰. The goal is to support cryptographic algorithms that are considered secure against an attack by a quantum computer. Currently this is not true for the most popular public-key algorithms, which can be broken by a sufficiently strong quantum computer.

HSMs can also be built to either address a range of application types or purpose built for specific use cases. As the name suggests, a general purpose HSM can support any application. It has a basic cryptographic command set that offers primitive cryptographic operations. Because its command set is primitive, developers that use such HSMs have to write the additional logic themselves. This increases development effort. Of course, the advantage is that the same HSM can be used in a wide range of applications that serve different industries and markets.

Cloud HSM

While HSMs offer a high level of security for cryptographic operation, they do come with overhead: they have to be procured, configured, and administered by a dedicated team. Organizations also have to forecast their needs ahead of time to allow for the latency of getting an HSM up and running. Cloud HSMs allow this burden of infrastructure overhead to be shifted to cloud service providers. All major CSPs have a Cloud HSM infrastructure that they maintain for internal cloud security and many of them also sell Cloud HSM services to their customers. Applications that use a Cloud HSM actually interface with it through a web service endpoint, and are completely unaware of the physical HSMs powering the service. This allows enterprises to adapt in real-time to changes in usage patterns. The same allure of elasticity and cost saving that has attracted countless enterprise applications to move to the cloud, has begun to appeal to the HSM market as well.

Secure Enclave

While HSMs are still the most trusted providers of RoT for key management and data protection, there are some emerging technologies that are vying for a piece of this cloud security market. One example is Software Guard eXtention (SGX)¹¹ from Intel that creates secure enclaves to execute code in a trusted and well protected silo. Applications running in these secure enclaves, and the data they manipulate at runtime, are protected against all external interference. This property can in turn be used to securely store and use cryptographic keys in CSP. Although this concept appears promising, security researchers and practitioners have found some vulnerabilities¹² which are being studied closely. As such, it may be some time before secure enclave technology is considered viable for mainstream key management systems and be on par with HSMs.

⁹ NIST Security Requirements for Cryptographic Modules, FIPS 140-2, <https://csrc.nist.gov/publications/detail/fips/140/2/final>

¹⁰ Post-Quantum Crypto Agility, <https://safenet.gemalto.com/data-encryption/post-quantum-crypto-agility>

¹¹ Victor Costan and Srinivas Devadas. "Intel SGX Explained." IACR Cryptology ePrint Archive 2016 (2016): 86. <https://eprint.iacr.org/2016/086.pdf>

¹² G. Chen. Et al., SGXPETRE: Stealing Intel Secrets from SGX Enclaves via Speculative Execution, <https://web.cse.ohio-state.edu/~zhang.834/papers/eurosp19.pdf>

Key Management

Security services, such as encryption and decryption of data, are based on cryptographic mechanisms that in turn rely on cryptographic keys. These keys have to be generated, distributed, stored, rotated, and eventually discarded in a secure and protected manner. Secure management of keys is one of the most critical elements when integrating cryptographic functions into a system. A security framework will be rendered ineffective, if the underlying key management is weak.

Good key management starts by having strong entropy, establishing a root of trust to store these keys, and then offering choices regarding how and where these keys are used. The scope of key management is broad, covering topics, such as key generation, storage, usage operations (e.g. encryption, decryption, signing, verification, etc.), lifecycle management (e.g. key rotation, backup, revocation, suspension, and deletion), along with logging and report generation capabilities. As keys control access to cryptographically protected data, key management becomes an important aspect of the overall security posture.

Federal Information Processing Standard (FIPS) 140-2¹³ is the benchmark for validating the effectiveness of cryptographic solutions. Although FIPS 140-2 is a U.S. Federal standard, compliance with it has been widely adopted around the world in both government and non-government sectors as a practical security benchmark and realistic best practice. There are four different levels (Levels 1-4) specified by FIPS 140-2. Depending on regulatory requirements, an organization can choose a key management system (KMS) meeting the relevant level. Good key management systems must be at least FIPS 140-2 Level 3 certified.

The adoption of cloud based data storage has led to an evolution of key management services to enable data encryption and decryption. Some noteworthy examples are: cloud-based native KMS, BYOK, HYOK, and cloud-based key brokering services. The sections below capture the main options available for such systems.

CSP Managed Key Management Service (KMS)

Cloud key management is the capability provided by cloud infrastructure providers to create and manage keys and control the use of encryption across a wide range of their cloud services as well as customer applications. Depending on the implementation by the cloud provider, Cloud KMS supports CSP managed keys along with customer managed keys. Main features and variations of these two options are listed below:

- **CSP Managed Encryption Keys in CSP KMS:** Some cloud providers, such as AWS, enable automatic key generation, when encryption is invoked by a service such as Amazon S3. These keys are referred to as CSP managed keys, and typically do not incur any added cost for the customer. While customers can monitor the key usage, they have little control over these keys.
- **Customer Managed Encryption Keys in CSP KMS:** Many cloud KMS offerings by CSPs support customer managed keys (CMEK). This capability allows customers to create keys, control access to the keys (leveraging the CSP's IAM solution) and manage the key lifecycle. The following options enhance security for customer managed keys in CSP KMS:
 - **CSP KMS with Hardware RoT (HSM):** Some CSPs provide this capability to offer key management services with hardware protected keys. This addresses compliance use cases and best practices that need keys to be protected in FIPS 140-2 Level 3 compliant key stores.
 - **CSP KMS with BYOK:** Most CSPs support customer generated key material importation into the CSP's cloud KMS. This key material can be generated by a customer managed on-premises solution (e.g. a HSM) or at a cloud-based key brokering vendor. An important differentiator in this approach is that the customer owns, operates, and controls the key creation. This helps customers with regulatory or internal requirements for key entropy.

Examples of various cloud key management and encryption options supported by AWS are listed in Table 3. While AWS is the dominant player in cloud and has certainly enjoyed first mover advantage, Microsoft Azure is also growing quickly and establishing itself as a leader offering cloud services for the Enterprise..

13 NIST Security Requirements for Cryptographic Modules, FIPS 140-2, <https://csrc.nist.gov/publications/detail/fips/140/2/final>

Table 3: Options offered by AWS based on target data for encryption

What is Encrypted?	How?	Where?	Storage	Comments
Volume or Disk Block	OS/Kernel	Compute instances e.g. AWS EC2	Local volumes (e.g. AWS EBS)	Entire disk volume is encrypted
Files or File Folders	OS/Kernel	Compute instances e.g. AWS EC2	Local file systems (e.g. AWS EFS) or external storage service (e.g. AWS S3)	Used for encrypting unstructured data
Application Data	Application	Compute instances e.g. AWS EC2	Local volumes (e.g. AWS EBS) or external storage service (e.g. AWS S3)	Used for encrypting structured data, such as a column of sensitive data in a database

Customer Managed Key Management Service (KMS)

For stronger key control, most cloud providers support customer supplied keys. This is also referred to as HYOK, as the keys are no longer stored in the CSP KMS, which offers stronger separation of duty. The keys are created and stored on the customer’s HSM, key brokering service, or key management device. With this method, the customer always maintains control of their key material instead of entrusting it to the cloud provider.

Customer managed KMS can also be used to enable encryption independent of the cloud infrastructure providers’ capabilities. For example, application-level tokenization and encryption of cloud data, such as Amazon Elastic Block Storage (EBS), Amazon S3 buckets, and Azure Files, etc. This is referred to as Bring Your Own Encryption (BYOE). Typically, a third-party KMS vendor’s solution is deployed as an on-premises or cloud-hosted appliance and managed by the customer. Some vendors also provide additional encryption tools to transparently protect data for certain cloud services and offer tools for extraction, transformation, and loading (ELT) of data. In addition to BYOE, third-party vendor KMS can also be used to support BYOK and HYOK use cases for cloud encryption.

Key Brokering

Key brokering refers to a third-party vendor solution used by an enterprise for cloud key management. Key brokers are available as cloud services or customer owned solutions independent from the cloud infrastructure provider. They enforce separation of duty between key management and data stores. Depending on the set of services provided by the vendor, key brokering could be used to enable BYOK, HYOK, or BYOE use cases.

Organizations use multiple CSP and hybrid clouds to meet requirements for different applications, to avoid lock-in, to maintain competitive pricing across vendors, and for high availability. Key brokering solutions also greatly simplify key management for multi-cloud and hybrid environments by creating consistent interfaces, centralizing management, automating key life-cycle tasks, such as key rotation, enabling audit, and logging support, and enforcing compliance requirements, such as FIPS 140-2.

Approaches for Cloud Data Protection

This section describes how principles of key management can be combined with security best practices to deliver key management and encryption solutions for the cloud. These solutions are illustrated in Figure 2 and represent a range of approaches with varying degrees of customer control. Each approach should be picked based on the value of data being protected and the level of control desired to achieve a given level of security and/or regulatory compliance. Attention should also be given to routine tasks that need automation and are difficult to achieve across different providers. If done incorrectly these tasks can lead to loss of compliance or service availability.

CSP Managed Encryption Keys		Customer Managed Encryption Keys (CMEK)		
Default CSP Protection	Managed CSP Protection	Bring Your Own Key (BYOK)	Hold Your Own Key (HYOK)	Bring Your Own Encryption (BYOE)
CSP transparently handles key management and encryption of data without any customer involvement or control	CSP handles key management and encryption of data; with some customer monitoring, but no control	Customer generated and rotates keys using entropy and security policies of choice and then hands these keys to the CSP	CSP handles encryption of data but the key management is under customer control; directly or through third party key broker	CSP only stores encrypted data. Encryption and key management is done by a customer provided solution or third party

Figure 2: Approaches for cloud data protection with varying degrees of customer control

The key management and encryption options, explained in sections below, are broadly classified into two categories: one in which the CSP manages the encryption keys, and the other in which this responsibility is handled by customers. The latter category of options is collectively referred to as Customer Managed Encryption Keys (CMEK) and gives customers greater control over protection of their data. BYOK, HYOK, and BYOE are specific ways of enforcing CMEK, each with a different level of customer control.

Default CSP Data Protection

In this approach the CSP transparently handles data protection for customers. This includes generation and management of keys, as well as use of these keys in encryption and decryption of the user data stored by the CSP. Many cloud providers offer this feature for their customers through mature solutions that can meet many low assurance data security compliance requirements. It can be an acceptable option for customers that either have relatively less critical data or do not want any additional responsibility related to protection of their data. However, they are choosing to fully trust their cloud provider. Different CSPs honor this trust in different ways. For example, some may deploy a strong root of trust to protect keys using HSMs, while others may not.

An example of default CSP protection can be AWS services, such as S3, that automatically create keys in AWS KMS for encrypting data stored in S3 buckets. The APIs that manipulate such data in S3 are completely unaware of how data is protected. Google transparently encrypts all data at rest.

Managed CSP Data Protection

The managed CSP protection approach is similar to the default CSP protection but offers customers some limited control over their keys. These keys are still generated and stored using a key management system provided by the CSP, but instead of handling all operations unilaterally and without customer input, the CSP exposes the key management interface to the customer. The customer therefore gets visibility into how their keys are generated, rotated, and used. Even though users do not fully control their keys, they can at least monitor them.

Bring Your Own Key (BYOK)

Some CSPs allow customers to generate their own keys and then import these keys into the KMS managed by the CSP. Customers can also opt to rotate these keys and provide the CSP with new versions. In this BYOK approach customers can enforce strong entropy and policy rules regarding key generation and rotation that may help meet regulatory compliance requirement. However, once these keys are handed to CSP, the key management and storage is done by the CSP. Most often the imported keys are used as key-encryption-keys (KEK). For greater control of key management, customers can use HYOK approach.

Hold Your Own Key (HYOK)

The hold-your-own-key (HYOK) approach offers the first real separation of duties between the CSP and the customer. In this approach the CSP still handles encryption and decryption of customer data but does not manage the keys. These keys are generated and managed by the customer, either directly or through an independent third-party, such as a key broker. The CSP can request access to these keys when encryption and decryption operations have to be performed. However, once these cryptographic operations are complete, and keys are no longer required, the CSP erases them from its cache. As such these keys are never allowed to persist by the CSP. The customer can choose to host the keys in a KMS in its own data center or use a separate cloud service.

Key management can be done in a cloud infrastructure or with a KMS hosted in the customer data center. Figure 3 illustrates HYOK with the key manager in the customer data center.

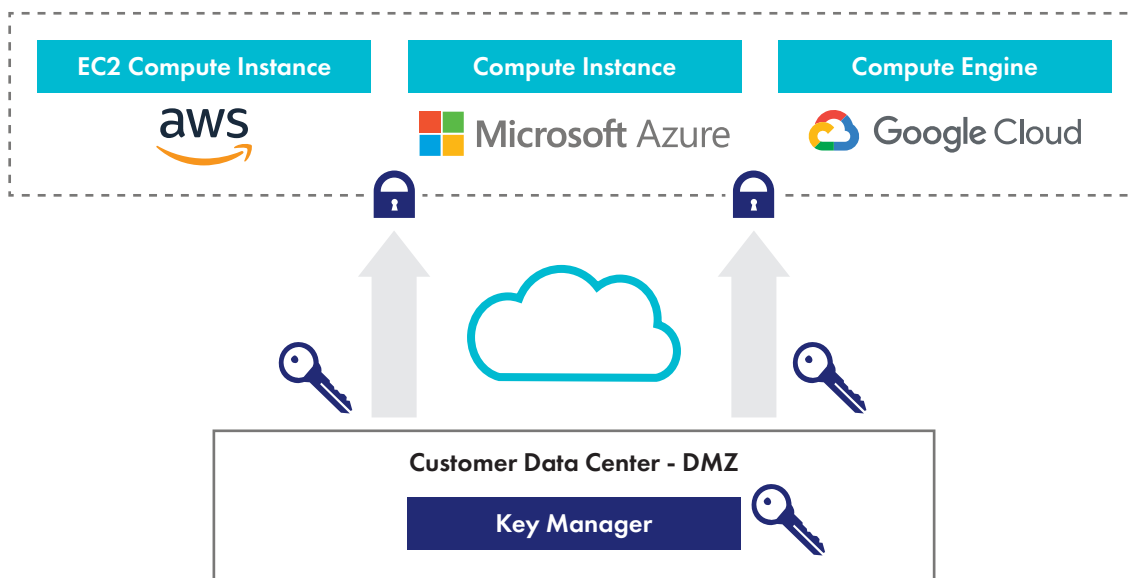


Figure 3: HYOK approach with KMS in customer data center and encryption done in a CSP

Bring Your Own Encryption (BYOE)

The bring-your-own-encryption (BYOE) approach offers the ultimate separation of duty by allowing customers to use their own encryption as well as key management tools instead of the corresponding solutions offered by the CSP. This gives customers the highest level of control over their data, since the data and keys are never exposed to the CSP. Instead, data is encrypted before sending it to the CSP for storage.

Conclusions

Recognizing the need for data protection, Cloud Service Providers, such as Google, Amazon Web Services, and Microsoft Azure, offer native encryption and Key Management Service with many of their offerings. Applications can use the encryption and KMS to protect data-in-use or data at rest. While the native encryption and key management services are robust, separation of trust between the CSP and the customer is not only a good security practice, it is often necessary for compliance. This separation of duties enables the customer to be the custodian of encryption keys and thereby achieve higher levels of assurance and risk management. Recognizing this important need for separation of trust, CSPs offer several approaches to allow a customer to be its own key custodian. These approaches are Bring Your Own Key (BYOK), Host your Own Key (HYOK) and Bring Your Own Encryption (BYOE). While BYOE is the best option for separation of trust between the customer and the CSP, BYOK and HYOK are good choices that meet many of the high assurance and regulatory compliance requirements. Figure 4 illustrates how the approaches described in this paper address separation of duties between CSP and customer for four independent operations: key monitoring, key generation, key management, and key use in cryptographic operations, such as data encryption/decryption. Key management implies complete key lifecycle management, including key rotation.

Approaches	Tasks							
	Key Monitor	Key Generation	Key Management	Data Encryption	Key Monitor	Key Generation	Key Management	Data Encryption
Default CSP Protection	✓	✓	✓	✓				
Manged CSP Protection		✓	✓	✓	✓			
BYOK			✓	✓	✓	✓		
HYOK				✓	✓	✓	✓	
BYOE					✓	✓	✓	✓
Responsible Party	CSP				Customer			

Figure 4: Approaches for separation of duties between CSP and Customer

This paper discussed:

- Encryption and key management principles
- The various flavors of encryption and key management services offered by AWS, Google Cloud and Microsoft Azure
- Multi-cloud key management
- Best practices for data protection in the cloud

The information presented in this paper should give readers a basic understanding of the various technologies and help them select the right approach to data protection in the cloud.

Acknowledgments

The authors would like to thank their Thales colleagues Charles Goldberg and Eric Wolff for their diligent review and very constructive feedback on this paper.

THALES

Contact us

For all office locations and contact information,
please visit cpl.thalesgroup.com/contact-us

> cpl.thalesgroup.com <

