Data Protection for Infrastructure as a Service

Enterprise Cloud Customer Council Technical Working Group Position Paper

December 2016

Preamble

Overview

The Enterprise Cloud Customer Council (E3C) is a group of enterprise companies, primarily financial services, with the vision to become a leading voice for enterprise organizations in the pursuit of cloud interoperability, security, and portability with major cloud industry vendors and standards bodies. The top cloud companies are closely involved and are actively assisting the E3C in achieving its goals. Intel Corporation serves as facilitator for the group.

Mission

The E3C mission is to develop requirements centered on common functional architecture, detailed specifications, and technical solutions that improve enterprise cloud adoption and integration while mitigating regulatory risk.

For the smooth and secure usage of hybrid cloud for applications and data, the council shall work toward

1) Enhancing security of clouds for secure network connectivity and data management that align to enterprise requirements
2) Ease of integration and life cycle management through interoperability across clouds
3) Ensuring enterprises have the same visibility and control of their resources in public cloud as on premises

Process

In order to achieve the E3C mission, the members form technical work groups focused on the top current impediments. The technical work groups work towards understanding common architectures and approaches and then establishing common requirements that form the E3C position papers. After the papers are approved by the E3C council members, consisting of executive level members from each of the member companies, the papers are published on a public website. Papers will be iterated accordingly as needed.

Usage of Position Papers

The position papers are not designed as prescriptive requirements but rather as input for organizations developing roadmaps and/or requirements relating to cloud adoption. These are freely available, however please reference E3C if you utilize these in RFPs, RFQs, publications, etc.

Document Conventions

RFC 2119 conventions apply to this document. For convenience, these definitions are provided here:

1. MUST: This word, or the terms "REQUIRED" or "SHALL", mean that the definition is an absolute requirement of the specification.

2. MUST NOT: This phrase, or the phrase "SHALL NOT", mean that the definition is an absolute prohibition of the specification.

3. SHOULD: This word, or the adjective "RECOMMENDED", mean that there may exist valid reasons in particular circumstances to ignore a particular item, but the full implications must be understood and carefully weighed before choosing a different course.

4. SHOULD NOT: This phrase, or the phrase "NOT RECOMMENDED" mean that there may exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.

5. MAY: This word, or the adjective "OPTIONAL", mean that an item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because the vendor feels that it enhances the product while another vendor may omit the same item. An implementation which does not include a particular option MUST be prepared to interoperate with another implementation which does include the option, though perhaps with reduced functionality. In the same vein an implementation which does include a particular option MUST be prepared to interoperate with another implementation which does not include the option (except, of course, for the feature the option provides.)

1. Overview

For enterprise customers, it is essential to maintain data integrity and confidentiality when leveraging public cloud services, as data is no longer within the customer's physical control. Customers also want the ability to verify that their data remains confidential and its exposure is minimized. The goal of this document is to provide guidance from enterprise customers to cloud service providers (CSP's) on methods of protecting data confidentiality via encryption for both data at rest and in transit.

Encryption is a piece of a broader security strategy. It provides assurance that if encrypted data accidentally falls into an attacker's hands, they cannot access the data without also having access to the encryption keys. With adequate implementation, if lower layers such as storage devices are compromised, data confidentiality can be maintained. Encryption also acts as access control and policy enforcement, centrally managed encryption keys create a single place where access to data is enforced and can be audited.

The scope of this paper is limited to Infrastructure-as-a-Service (IaaS) offerings and is intended to facilitate further discussion with cloud service providers on data protection. There are several other important security considerations, such as application security, identity and access management, configuration management, datacenter security, business continuity, etc., all of which should be included in customer's security strategy but are outside of the scope of this paper.

2. Data at Rest: Integrity & Confidentiality

Customer data may include both content directly and indirectly generated by the customer. Content directly generated or supplied by a customer includes, but is not limited to, data stored in object storage, block storage disks/snapshots, and operating system images. Content indirectly generated by a customer may include data generated by a CSP on behalf of a customer, including logs, performance metrics, and resource metadata.

This paper will discuss data protection guidance with respect to two levels of data classification or sensitivity, defined below:

1. Normal Sensitivity: Customers expect data to be protected, but with a security tradeoff where a CSP generates/supplies encryption keys and manages the encryption key lifecycle end-to-end. This may be a reasonable approach to balance performance, operational risk, and operational complexity. Examples of data that may be classified with normal sensitivity may include indirect customer data (CSP-generated log, metrics, metadata, etc.), publically available datasets, and dev/test data. It would explicitly exclude data classified as Personally Identifiable Information (PII) or Material Non-Public Information (MNPI).

2. High Sensitivity: Customers expect to play a role in managing the key lifecycle, either by explicitly managing or supplying encryption key material. The expectation is that although keys may be exposed to the provider services during cryptographic operations, if architected correctly the risk of key exposure can be minimized.

   The customer managed model is a shared responsibility between CSP and the customer. Customers understand their own maintenance responsibilities and keep track of key materials and its use. CSP access to customer managed keys must require customer interaction; this may be asynchronous operation such as granting permissions to specific service and access model, or importing key materials.

   As reference architecture, CSP may extend their centralized Key Management Service to offer mechanisms for enterprise customers to provision, manage, expire, revoke, delete, and rotate their own keys. This may involve a Hardware Security Module (HSM) located in CSP infrastructure but managed by the customer, in a customer data center, or at another CSP. Regardless of implementation, customer must have exclusive access to encryption keys.

3. Data at Rest: Key Management Lifecycle

For both data sensitivity classifications, customers expect CSP's to implement a Key Management Service (KMS) to manage the encryption keys that protect customer data. CSP's may have different offerings or flavors of a KMS, but a KMS itself must be an integrated service for generating, distributing, and managing cryptographic keys for resources and applications.

Below are some definitions used in the remainder of this document to further describe a shared responsibility Key Management System.

Term	Acronym	Definition
Key Management Service	KMS	An integrated service for generating, distributing, and managing cryptographic keys for resources and applications.
Key Encrypting Key	KEK	A key used to protect other keys.
Data Encrypting Key	DEK	A key used to encrypt content.
Hardware Security Module	HSM	A physical computing device that safeguards and manages digital keys for strong authentication and provides crypto processing.
DEK wrap	N/A	The act of encrypting/encapsulating a DEK with a KEK.
DEK unwrap	N/A	The act of decrypting/un-encapsulating a KEK protected DEK.
DEK caching window	N/A	The lifetime for which a DEK may be cached in memory.

4. Data at Rest: Key Encryption Keys

Key Encryption Keys (KEK), should be protected throughout their lifecycle and be maintain in a key management service. Loss of control of KEKs may result in clear text exposure of data encryption keys, thereby affording attackers who may already have access to customer's encrypted data the potential to decrypt the data.

For Normal and High Data Sensitivity Classifications

1. KEKs must not be moved or stored in plaintext.
2. The CSP should provide the ability to require approvals and/or a waiting period before key deletion to avoid accidental deletion.
3. CSP KMS must provide a mechanism to securely rotate, expire, revoke, and delete KEK's.
   a) Normal Data Sensitivity: May be implemented by a CSP's own policy and schedule.
   b) High Data Sensitivity: CSP must provide the customer the ability to conduct all key operations listed, with the additional operation of importing a KEK.
4. CSP KMS should provide access control capabilities to define key scope and usage policies.
   a) Normal Data Sensitivity: May be implemented by CSP defined policies, where scoping is applied on per service and/or per-customer account basis.
   b) High Data Sensitivity: CSP must provide customer ability to restrict key usage to only authorized services, resources, and/or users.

For High Data Sensitivity Classifications

5. A KEK must be able to be generated and persisted on a customer-managed Hardware Security Module (HSM). This HSM may be hosted on a customer's premise or by the CSP.
6. If the HSM responsible for the KEK is hosted by the CSP, to ensure customers maintain control of the KEK, HSM's hosted by a CSP must provide for the following:
   a) The CSP administrators must not have privileged access such as security officer or partition (or equivalent) access.
   b) The HSM must provide a dedicated logical or physical partition to a single customer.
7. For compliance purposes, some customers may require that a CSP hosted HSM meet or exceed the requirements defined by FIPS 140-2 Level 2+.
   a) Some customers require HSM operation, compliance, and/or validation at FIPS 140-2 Level 2+. The CSP should provide a version of an HSM that meets these standards.

5. Data at Rest: Data Encryption Keys

Data Encryption Keys (DEK) should be protected throughout their use and not expose its plaintext unnecessarily. Exposure of data encryption keys will result in data loss event if an attacker already has access to customer's encrypted data.

For Normal and High Data Sensitivity Classifications

1. A DEK may be generated by the CSP provided all DEKs are cryptographically protected by a KEK immediately after creation (DEK wrap).

2. DEKs can only be written to non-volatile storage in a KEK protected state.

3. DEKs not protected by a KEK may only be stored in volatile memory and must not be persisted to non-volatile storage via direct action or memory management events. Hardware-based secure enclaves (e.g. TPM, SGX) should be used where possible and when prevalent to secure DEKs in memory. DEKs must be securely flushed from volatile memory after expiration.

For High Data Sensitivity Classifications

4. DEKs may be protected by multiple KEKs, or a hierarchy of KEKs, so long as the customer-managed or supplied KEK is always required to unwrap the DEK.

5. If a customer disables, expires, revokes a KEK from the CSP, DEK unwrap requests requiring the KEK must fail and be logged.

6. All instances receiving unwrapped DEKs must provide audit logs for the following:
   a) Unwrapped DEK issuance to the node
   b) DEK renewal and expiration/flush events with correlation to the originating DEK issuance event

7. Instances receiving unwrapped DEKs should not delegate or distribute DEKs to other instances.
   a) During compute instance live migration events, any DEKs issued to the source hypervisor should not be copied to the target hypervisor as part of the running memory transfer. Instead, the target hypervisor should explicitly request and receive an explicit grant to the encryption keys required by the instance. This is to ensure an audit trail and proper chain of custody of the encryption keys is maintained.

6. Data at Rest: Key Usage Auditing

For Normal and High Data Sensitivity Classifications

1. Provider must maintain an audit log of all events involving a customer's keys within the CSP's Key Management System. This includes but is not limited to:
   a) KEK creation, import, rotation, expiration, revocation, and deletion events.
   b) DEK creation, wrap, unwrap, revocation, and expiration events, with a reference to any associated resource(s) if applicable.

2. CSP KMS audit logs must be accessible to the customer on a near real-time basis, and be retained in an immutable state for a defined or configurable period of time.

For High Data Sensitivity Classifications

3. From the KMS auditing events, the customer must be able to monitor and analyze for the following event patterns:
   a) Every request to or within the CSP that requires usage of a customer's KEK should have a corresponding DEK unwrap event.
   b) If the DEK was issued/transferred to a resource, the above-mentioned DEK unwrap request should have a corresponding DEK expiration/flush event.

7. Data in Transit Integrity & Confidentiality

While modern applications may be designed with security in mind from the outset, many customers may migrate legacy applications to CSP's that rely on insecure network transports. Providing a method of ensuring confidentiality and integrity of data in transit can help ease the migration of such applications to a CSP environment.

1. The CSP should provide the customer the ability to construct an encrypted network transport at the infrastructure layer (i.e. layer 3), such that:
   a) Network transit between a customer's instances within in the CSP are encrypted.
   b) Data movement between the customer site and the CSP are encrypted.
   c) The encrypted transport should not be exposed by a single shared secret.

2. The CSP may provide the customer the ability to define authentication credentials for resources over the network (e.g. unique X.509 certificate for each compute instance).

8. Data in Transit Visibility & Transparency

Many large regulated customers use a number of tools beyond firewalls to analyze and mitigate network-based attacks, including denial of service, intrusion detection/prevention, and traffic recording/forensics.

1. The CSP must maintain an audit log of all changes made to a customer's network configuration.

2. The CSP must provide the ability to log network flow metadata (e.g. IPFIX) for all inter-resource communication on the CSP virtual network. The CSP should provide a facility to do partial or full network packet capture (e.g. pcap) against instances.

3. The CSP should provide the capabilities to constrain and verify network configuration and analyze network flow to validate expected security state versus intent, in near real-time.
   a) Examples of configurations and intents a customer may want to verify:
      - A given network should not be routable to the public Internet
      - A given instance should not be allowed to have both a public IP address and have connectivity (e.g. route) to the customer's premise

4. The CSP should provide the customer capabilities to detect and alert on anomalous network traffic in near real-time (e.g. IDS/IPS capabilities).

5. The CSP should have a documented process for handling and mitigating DDoS attacks. If a DDoS attack is targeting a specific customer, the CSP should have a process for notifying and sharing relevant information about the event to the customer.