Security Specification for Cloud Data Services

Enterprise Cloud Customer Council Technical Working Group

October 2017

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

This document outlines specification for security control relevant to protecting customer data when using cloud based Data Services. The scope of Data Services includes databases, object storage, messaging queues, shared file systems, and streaming data processes. The paper expresses customer position on authentication, authorization, data protection for at rest and transit, and transparency requirements in order to prevent the exposure or modification of the customer's data

2. Definitions

Some definitions used in the remainder of this document, to further describe a shared responsibility Key Management System.

Term	**Acronym**	**Definition**
Cloud Service Providers	CSP	Cloud Service Providers
Key Management System	KMS	An integrated approach for generating, distributing, and managing cryptographic keys for devices and applications.
Customer Managed Key	CMK	Customer managed encryption key stored within Cloud Service Provider's KMS. The implementation must allow customer to control the lifecycle and define usage of the key
Identity and Access Management	IAM	A security discipline and system implementation that addresses the need to ensure appropriate access to resources

3. Identity & Access Management

3.1 Federated identity

1. CSP must support open standard for identity federation and delegated access, including but not limited to OpenID Connect (OIDC)/OAUTH 2.0 and Security Assertion Markup Language (SAML) 2.0

2. The CSP should support any customer defined identity provider (idP), as long as the external idP is compatible with the agreed standard for identity federation

3. Federated users should be granted access without requiring customers to create a separate IAM user in CSP authentication system

4. CSP federation service must not have a dependency on customer internal directory service, or requiring customer to export their user data from customer internal directory service

5. All APIs call shall support federated identity

3.2 Granular controls for managing/restricting access to services

CSP's Identity and Access Management (IAM) must grant the individual or the service with just-in-time access to the right resources, with just-enough access, and for the right reason

1. CSP must provide a permission policy that can be associated with an identity (such as a user, a group, or a role), a resource (such as a message queue), or both. The permission policy must allow customers to define fine-grain operations of their individual users, such as control over the permission to provision, modify, manage and delete their Data Service resources

2. The permission policy must include elements for customer to define
    a. Identity/role principal
    b. resource operations - actions that the user can perform
    c. resource - this should allow user to specify granularity to the lowest object level possible for the resource
    d. conditions when a policy should take effect. Conditions must include but not limited to following:
        1. IP/CIDR of the request matched customer defined allowed sets
        2. Multi-factor authentication protected API access
        3. Date/time of allowed access

3. CSP shall provide an implementation of Role Based Access Control (RBAC). Permissions must be of additive model, where default is "deny all". Permission added only if customer explicitly declared within a permission policy

4. If multiple policies are associated, where a user or a service has multiple roles, the permission shall be a union of policies attached to each role

5. CSP must provide customer ability to report on and evaluate the security of the resulting composite policies

6. CSP should provide customer ability to govern IAM at top-level/organizational level; allowing customer to manage cross accounts and projects

4. Network Access

4.1 Private Network Access

1. Data Service offering, where possible, shall allow the customer to host within customer own virtual private cloud (VPC) / virtual network

2. VPC security controls shall be available for customer to leverage and apply to the Data Service (e.g. virtual firewall to control network access to instances running within the data warehouse clusters)

3. For fully managed Data Services, CSP shall provide private routing access for customer VPC to Data Service directly
    a. Access must not require routing via the Internet
    b. Access should be routable from on premise network without routing via internet or proxy infrastructure

4.2 Transport Security

1. Data Services must provide the ability to encrypt at the transport layer for customer and the application to connect to the Data Service, e.g. TLS/SSL connections to a data warehouse cluster

2. For TLS/SSL, Data Service should provide the option to support both CSP generated certificate, or customer generated/managed certificate. CSP must apply standard certificate validation (validity, chain, expiry, revocation)

3. Certification pinning capability should be available in CSP client SDK. For example, some CSP java client SDK currently trusts default client cert store. SDK should be able to perform cert pinning based on CSP issuing CA

4. All system connections between Data Services must also support encryption in transit by default, e.g. communication between data warehouse cluster and the object storage, when conducting data loading or backup operations

5. CSP must document explicitly where data in transit is traversing in clear over CSP network

5. Data at Rest & Key Management

Data Service must provide capability to encrypt all data stored on disks, file system, within the cluster, intermediary data that persists, and all backups within the Data Service eco-system. Where possible, both server-side and client-side encryption strategy should be employed

5.1 Encryption

1. CSP must provide Customer Managed Keys (CMK) offering within its own KMS, and Data Service integration adoption to use KMS. Implementation should adhere to NIST published guidelines and specification described in the IaaS data security paper

2. CSP Data Services must clearly define and document encryption keys usage, including frequency, calling identity, and handling of keys. Data Services must log all CMK usage

3. Based on required usage explicitly defined by Data Services, customer should be able to formulate policies to control key scope, policies should be expressed in following terms
    a. Identity/Data Service principal
    b. Resource – as granular as possible. E.g. a specific column or cell in a database table, a specific topic in message queue
    c. Usage scope or conditions -e.g. Grant usage to Data Service for indexing purpose, time of day

4. Data Service requesting key must adhere usage under condition it is granted. Key material should be purged from memory or cache after usage, and should not be available for usage outside of explicitly granted condition

5. Data Services should provide support for server-side encryption with CMK, client-side encryption with CMK, or both, depending on usage pattern. Following table provide high-level guidelines and examples:

	Usage Pattern	
	CSP service introspect Customer Data	**Platforms usage only**
Pub / Sub Services (Queues, Streams)	Server-Side Encryption with CMK. An example is a message queue where the CSP service is a publisher, consumer, or both. CSP service must encrypt data/payload for PUT request into the queue, and decrypt data for authorized subscriber IFF subscriber has access to CMK for GET request from the queue. CSP may leave metadata or header unencrypted, such as preserving FIFO queues.	Client-Side Encryption with CMK. This example is where CSP service is not involved in consuming or publishing, where CSP service has no reason to introspect data.
Storage (object store)	Server-Side Encryption with CMK, or Client-Side Encryption with CMK. An example is object storage used for ETL or backup/recovery within CSP Data Service eco-system.	Client-Side Encryption with CMK. This example is where customer application is using object storage purely for storage purposes.

Data Warehouse	Server-Side Encryption with CMK. Furthermore, CSP should move toward client-side encryption with CMK. Explore Homomorphic encryption schemes and perform calculations on encrypted data without decryption.	NA

5.2 Implementation transparency

1. CSP must document explicitly where and how is encryption implemented. For example, for database encryption, whether its columnar encryption, cell level encryption, cluster level, etc.

2. CSP must adopt reasonable encryption standards, with its implementation tested and verified by third-party with subject matter expertise

3. CSP must document explicitly where encryption for data-at-rest is not supported, anywhere along data path where customer data is processed within the Data Service eco-system. E.g. loading from object store, Database not supporting columnar encryption, indices, etc.
    a. All Data Service key requests must be logged, with request details. Logs must accessible to customer
    b. Any usage outside of expected use are considered as unauthorized, therefore, most notify customer
    c. Intermediate key-cache – customer should be able to specify time duration of cache, request count of the cache
    d. CSP should implement anomaly detection for encryption key usage to identify potential unauthorized/unexpected usage patterns

5.3 Practical implementations of secure multi-party computation and homomorphic encryption

1. CSP should engage in crypto advancement, be involved in research, and implement client-side encryption and process encrypted query where practical

   For Example:
   Usage of homomorphic encryption to process encrypted data, encrypted queries, within data warehouse services. Secure multi-party computation for calculation or encryption key sharing

5.4 Auditable Data Destruction

1. CSP should provide mechanisms such that customers can reliably and consistently designate data for destruction/removal. The mechanism should comply with the audit/tracking requirements for the assorted regulations that Financial Institutions are tasked with meeting.

2. Data Destruction mechanisms should be available for both block and object storage as well as platform like systems such as databases or HSMs as a service.
3. As customers have varying requirements for what is acceptable for data destruction, CSP shall provide a "targeted data destruction" mechanism where, for example, guarantees are made with regards to data overwrites. Additionally, CSP shall provide a "crypto shedding" option where destruction of crypt key material alone is guaranteed.

6. Service Provider Administrator Access

1. All CSP administrative access must be just-in-time, least privilege for non-routine production access

2. CSP administrators must not have direct access to unencrypted customer data. Administrator access to customer data must require customer approval, and access must be removed immediately after work has been completed

3. All CSP admin activities must be logged. CSP must implement an internal review process for its usage of administrative/privileged access

7. Auditing & Visibility

1. CSP must implement logging of customer-requested administrative actions

   For example:
   On customer request, CSP administrator may need access to the guest OS to troubleshoot a data warehouse instance. The access and its activities must be logged and monitored, when access is no longer required, access privileged must be revoked.

2. CSP must implement logging of all data access/query events

3. All CRUD operations conducted by Data Service processes, e.g. Create Table, Drop Users, Submit Jobs, Get Object Storage, Copy, Load, etc.

4. Data Service must log each query before it is run on the database / data warehouse. The log operation should not block the query execution

5. Data Service offering must provide customer visibility into performance and transaction metrics

6. CSP must implement logging on all API/CLI/Console access

7. Logging must capture all identities; users, application, Data Service processes and its connection events (authentication attempts, connections, disconnections)

8. Logs must be made available and exportable to customer.

a. Logs must be immutable
b. There shall be a documented log guarantees
c. If a security-sensitive service such as IAM had logging failures, and failed to log changes or access to its service, or that logging services had general failures or data leakage events, CSP must conduct root-cause analysis on impacted changes and inform customer in a timely manner
d. Log retention shall be configurable by customer

9. CSP must provide notification of provider-made changes or maintenance that may impact the customer

8. Change Management and Release Controls

1. CSP must implement controlled and phased release of updates. Where patches and upgrades are applied during a configurable maintenance windows

2. CSP must provide release notes and announcements, security review details

3. CSP must implement rollback policies in the event of functionality or performance regressions

9. Data Domicile and Residency

1. CSP implementation must provide evidence to assure customer and regulatory body that data domicile is strictly enforced. Where customer data being placed in one region, is not stored, backed up, or persisted in any means in another region

2. CSP must provide customer capabilities to constrain user access to a specific region, or must deny user access if region requirements are not met

10. Security Program Management

1. CSP is responsible for protecting its infrastructure and services offered. CSP must adhered to standards (e.g. ISO 27001) and industry best practices for its own security management practice

2. CSP must follow Secure Software Development Lifecycle best practices. This includes but not limited to formal design review, threat model, code review, risk assessment, and penetration testing

3. CSP must demonstrate how it achieves key compliance controls (e.g. SOC 1-3), regulatory requirements and regional laws (e.g. GDPR). CSP must provide artefacts such as examination reports or attestation via independent 3rd parties

4. CSP must document explicitly which of its products/services are compliant, and products/services that are non-compliant to key compliance controls, regulatory requirements and regional laws