White Paper

Cisco Context Service Security and Privacy

First Published: February 8, 2017

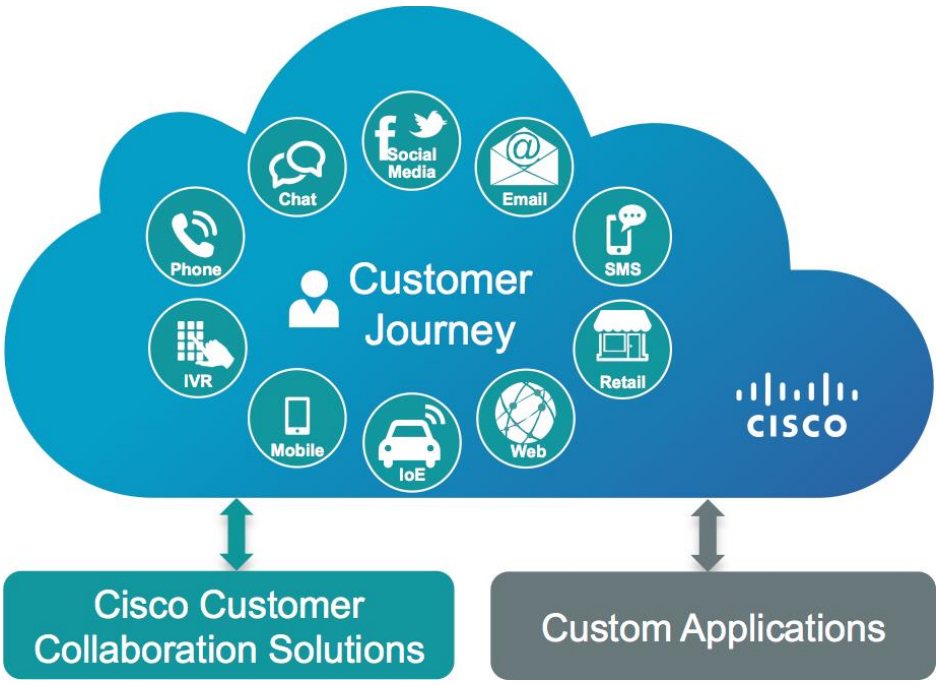

Cisco Context Service is a secure, cloud-based storage service that provides a repository for customer journey data.

This white paper provides a security and privacy overview of Context Service and the integration points for applications that use it.

The Cisco products, services, and features described in this document are in varying stages of development. While some of such Cisco products, services, and features currently exist, others are under development or planned for the future. Find additional information at www.cisco.com. Cisco will have no liability for delay in the delivery or failure to deliver the products, services, or features set forth in this document.

Table of Contents

Table of Figures

Overview

What is Context Service?

Context Service is a secure cloud-based storage service that provides a repository for customer journey data. It enables Cisco Contact Center customers to deliver a seamless omnichannel experience with an out-of-the-box integration from Cisco Customer Collaboration products. It contains simple integration points, allowing third-party applications to create, store, update, and retrieve data.

- Context Service provides a flexible data store for storing customer interaction data. Businesses can define what customer interaction data they want to store and how to store it.
- Cisco's Contact Center products (Unified Contact Center Enterprise, Packaged Contact Center Enterprise, Unified Contact Center Express, and Hosted Collaboration Solution-Contact Center) come with an out-of-the-box integration with Context Service. Context Service comes with an API that integrates with front-end, back-end, retail, or Internet of Things (IoT) applications to capture a complete view of the customer journey.
- Context Service delivers SDKs. These SDKs wrap the business logic of the client-side encryption model and manage the interactions with the services involved (such as Cisco Common Identity or Cisco Spark KMS). Cisco products use the same SDKs that are released.
- The service provides a secure, easy-to-use interface that allows customers to fully manage their own security. For information on why this is important, see the section "Security and Privacy Challenges for Cloud Collaboration" in the Cisco Spark Security and Privacy White Paper.

Context Service Encryption and Data Security

Context Service integrates with Cisco Common Identity, a federated identity provider that supports Single Sign-On (SSO) with customer organizations.

Context Service uses a multipronged security model. First, like most cloud-based services, Context Service enforces access control at the application level. Users can only access data within their own organization and workgroups.

Second, Context Service re-enforces the access control model with client-side encryption. This encryption allows the organization (not the cloud provider) to fully control access to the data. To read the data, you must have access to the keys to decrypt it. This is performed with strict separation of the Key Management Servers from the Cloud Services.

Spark Security

Context Service uses the same set of Key Management Servers (KMS) as Cisco Spark Messaging.

For more information about how KMS works, what it means for your enterprise and organization, and the features it provides, see the Cisco Spark Security and Privacy White Paper.

Context Service Data Types

Context Service follows a tiered data encryption model. The concept is simple. Let each organization define the sensitivity of its data. Data is divided into three categories:

- Personally Identifiable Information (PII): Data that contains personally identifiable information (employees, customers, contacts). Examples include names, email addresses, or phone numbers. This data is encrypted at the client without giving the cloud service access to the keys.
- Encrypted (Non-PII): Data that is considered sensitive to the company or organization. This data may include trade secrets or business leads. If this data is compromised, it could result in loss of intellectual property or harm to the business. Examples include notes, problem descriptions, or subject or summary of issues. This data is also encrypted at the client without giving the cloud service access to the keys.
- Unencrypted: Data that is not sensitive to the company, does not contain personal information, and would not result in loss of intellectual property if read by a third party. Examples include dates, tags, or reference URIs. This data is not encrypted but is still protected through access control measures. Keeping certain types of data unencrypted allows for different methods for searching it. These methods include (but are not limited to) ranges, greater than and less than operations, substrings, and computations (such as average, min, max).

By separating personally identifiable data from other sensitive data, the business can share subsets of its data with other groups within the business or with other businesses.

How It Works

Objects and Associations

Context Service consists of eight separate objects:

- Organization: At the top is the organization. The organization represents Cisco's customer and has a direct mapping to an organization in Common Identity. An organization may be an instance of Unified CCE, Packaged CCE, Unified CCX, or HCS-CC.
- User: A user is part of an organization. Users can be machine accounts (used to represent non-human processes, such as an IVR system) or actual users (authenticated using SSO).
- Workgroup: A workgroup represents a subset of an organization. Workgroups are used to control access to runtime objects. Users belong to one or more workgroups. POD, Customer, and Request objects are associated with one or more workgroups that can access them.
- Field: A field is a definition of an element. It contains classification (PII, Encrypted, Unencrypted), data type (String, Integer, Double, Boolean), and Localization. Access is controlled at the organization level. Cisco provides base sets that are accessible for all organizations. For example, Context_First_Name is classification PII with data type String. For more information, see the Cisco Context Service SDK Guide.

- Fieldset: A fieldset is a logical grouping of a set of fields. Access is controlled at the organization level. Cisco provides base fieldsets that are accessible to all organizations. For more information, see the Cisco Context Service SDK Guide.
- POD: Piece of Data. A POD represents a single interaction in the customer journey. It contains data (PII, Encrypted, or Unencrypted) surrounding a distinct interaction with an end customer. This data contains attributes (State, Media Type, Fields and Fieldsets) specific to the business. A POD may be associated with 1 or more workgroups, 0 or 1 customer, and 0 or 1 request.
- Customer: Contains data (PII, Encrypted, or Unencrypted) about an end customer (such as phone number or email). This data contains fields and fieldsets specific to the business. A customer may be associated with 1 or more workgroups and 1 or more PODs.
- Request: Contains data (PII, Encrypted, or Unencrypted) about a customer request. This data contains fields and fieldsets specific to the business. A request may be associated with 1 or more workgroups and with multiple related PODs.

Figure 1: Context Service Objects and Associations

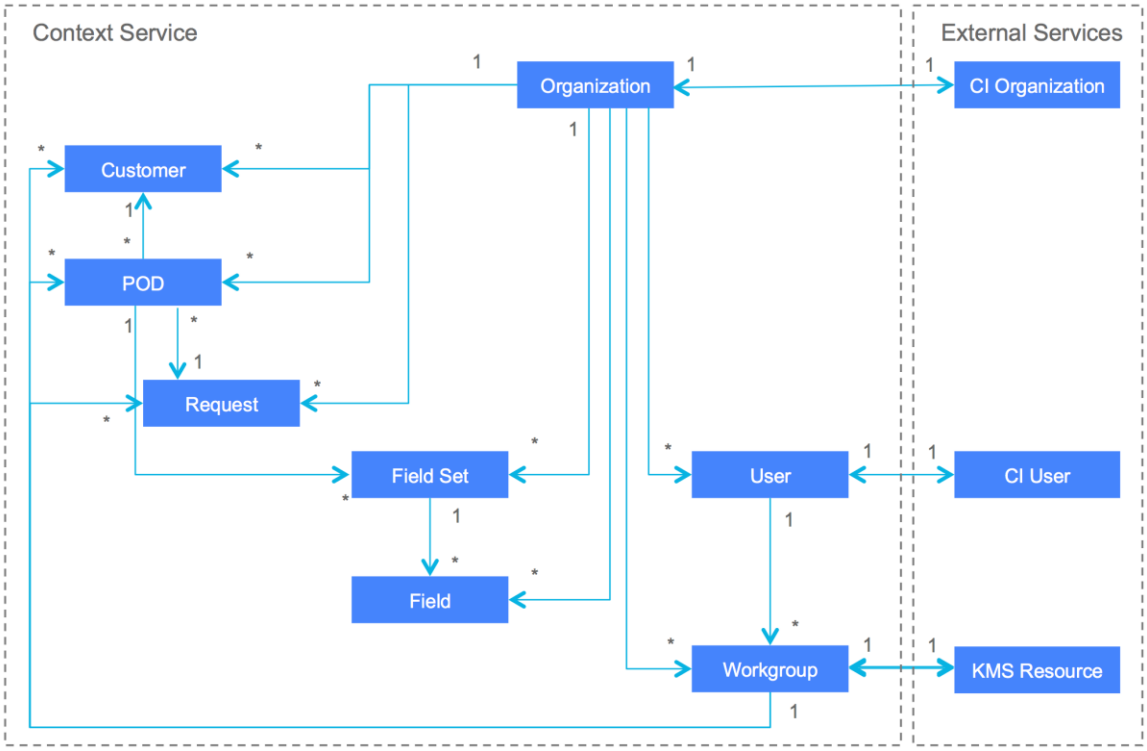

Workgroups

As previously mentioned, Context Service re-enforces the access control model with client-side encryption. Maintaining strict separation of the KMS from the Cloud Services means that the organization controls key access.

The Cisco Spark KMS has three main objects: Resources, Keys, and Users.

- Keys are bound to resources and used to encrypt data.
- Users are authorized to resources.
- Users can retrieve any keys bound to authorized resources.

In Context Service, there is a one-to-one mapping between workgroups and KMS resources. Currently, Context Service supports Production and Lab modes or workgroups.

Encryption and Hashing

Context Service applies an encryption scheme to secure data and a separate hashing scheme to enable lookup of that data. First, each set of data within an object is encrypted with its own unique key. If the object contains both PII and sensitive non-PII data, separate keys are used to encrypt each set of data. These keys are generated with the Context Service SDK.

To encrypt, store, and control access to this data, Context Service uses a concept known as a Secure Content Resource (SCR).

An SCR contains the location, algorithm, and keys needed to access the data. An SCR is created for each data section (PII and Encrypted) for each object. The SCR is then encrypted with a key for each workgroup that has access to the object. The decrypted SCR is the same for all workgroups. However, it is stored in an encrypted format within each object separately for each workgroup, along with a pointer to the key to decrypt the SCR.

Figure 2: Example of a Secure Content Resource (SCR)

Secure Content Resource

```
{
    "enc": "A256GCM",
    "key": "ZMpktzGq1g6_r4fKVdnx9OaYr4HjxPjIs7l7SwAsgsg",
    "loc": “pod.piiData”,
}
```

Due to the nature of the encryption scheme (each object is encrypted with its own keys), searching for data can be problematic. To allow the lookup of PII and other sensitive data, a hashing technique is used.

A special KMS key is created when you enable Context Service that is bound to a resource for the organization. Every user in the organization is authorized to this resource.

Each time that an object is created or updated, a SHA-256 HMAC, using the special KMS key bound to the organization as a salt, is used to compute a one-way hash over each element. These hashes are then stored alongside the data in Context Service and indexed to allow lookup of the data.

Additional Notes:

- Each element is converted to a lower case string before hashing. This provides case-insensitive search for all languages except Lithuanian (lt), Turkish (tr), and Azeri (az).

- This scheme provides for an exact match lookup only. Other techniques are being investigated that would allow for partial match using wildcarding.
- Unencrypted data is also hashed, but with a generic, system-wide key rather than the organization salt. This is needed to support searches where the field type is not known and is used to keep from unintentionally passing sensitive data to the web applications.

Enabling Context Service

When a system enables Context Service, the following operations are performed:

First time only:

- Resources for Org, Lab, and Production workgroups are created in KMS and their IDs are registered with Context Service.
- KMS for HMAC Salt is created, bound to the Org resource, and Key ID is registered with Context Service.

All:

- Machine accounts for Lab and Production workgroups are created.
- Machine account IDs are registered as users of Context Service for the respective workgroups.
- Machine accounts are authorized to the KMS resource that matches the workgroup it is associated with.
- Credentials and URLs are encoded into a string and returned to the user.
  IMPORTANT: This data encompasses passwords and must be protected.

Runtime Operations

Context Service tracks customer interactions using POD, Customer, and Request objects. To simplify the interface for users (including Cisco Collaboration products), we publish SDKs that internally wrap the business logic of interacting with Common Identity (Authentication), KMS (Key Management), and the Context Service APIs.

To simplify usage, objects exposed through the SDKs do not refer to the classification of the data used. This is handled by building out fields and fieldsets in the Context Service Dictionary during configuration.

Figure 3: Internals of the Context Service SDK - Bucketing, Hashing, and Encrypting

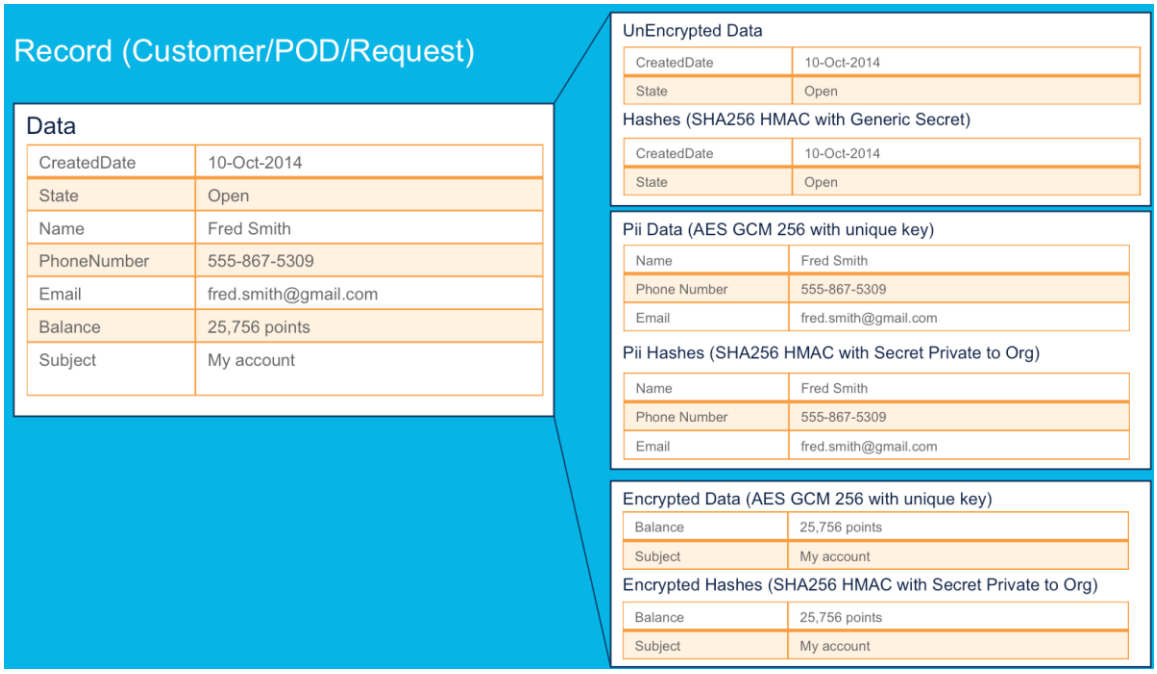

The SDKs expose the following methods: Create, Get, Update, Search, Delete (Lab only)

The following sections describe the interworking of the SDK when each of these methods is called.

Create/Update

This diagram shows a typical create or update call. It contains the computations performed by the SDK and the interactions with the cloud services. The SDK employs caching when possible to reduce the number of calls to external services, which reduces the overall execution time.

Figure 4: Context Service SDK - Create/Update Business Logic

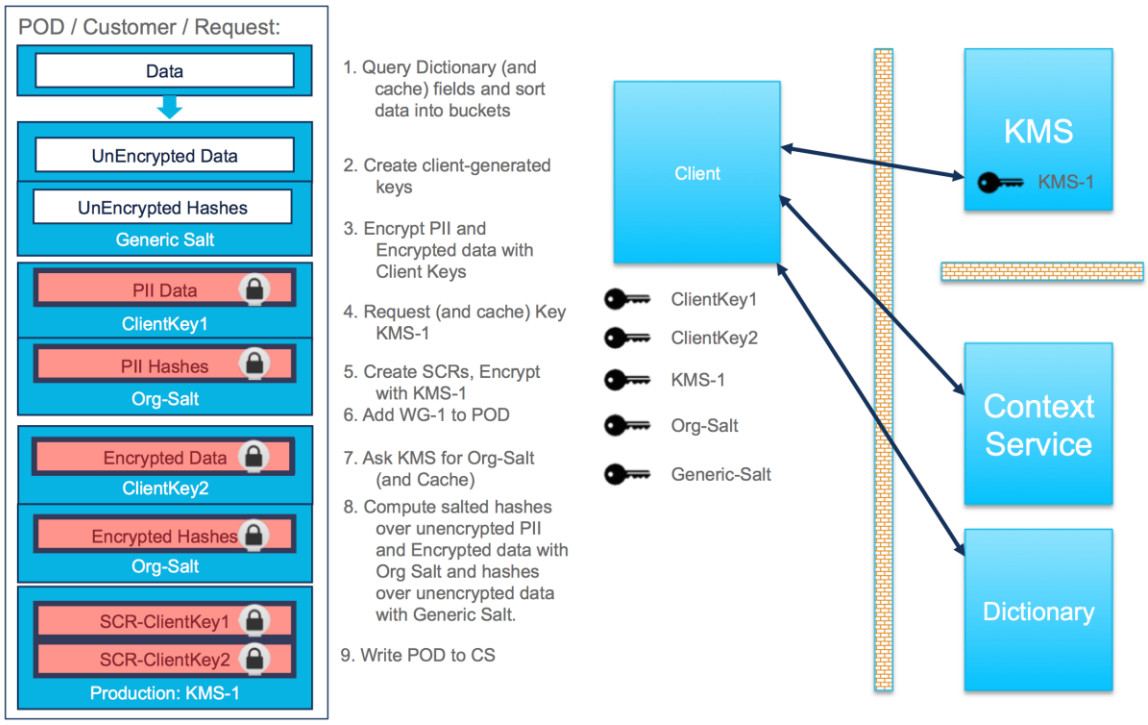

Get

Figure 5: Context Service SDK - Get Business Logic

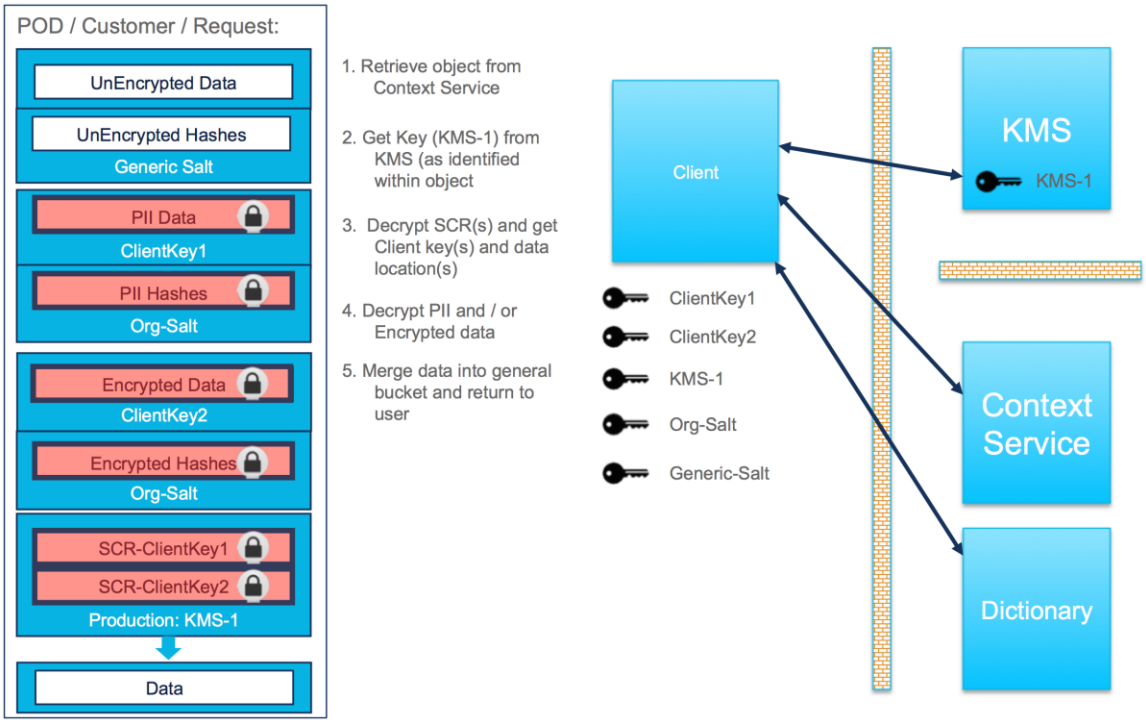

Search

Figure 6: Context Service SDK - Search Business Logic

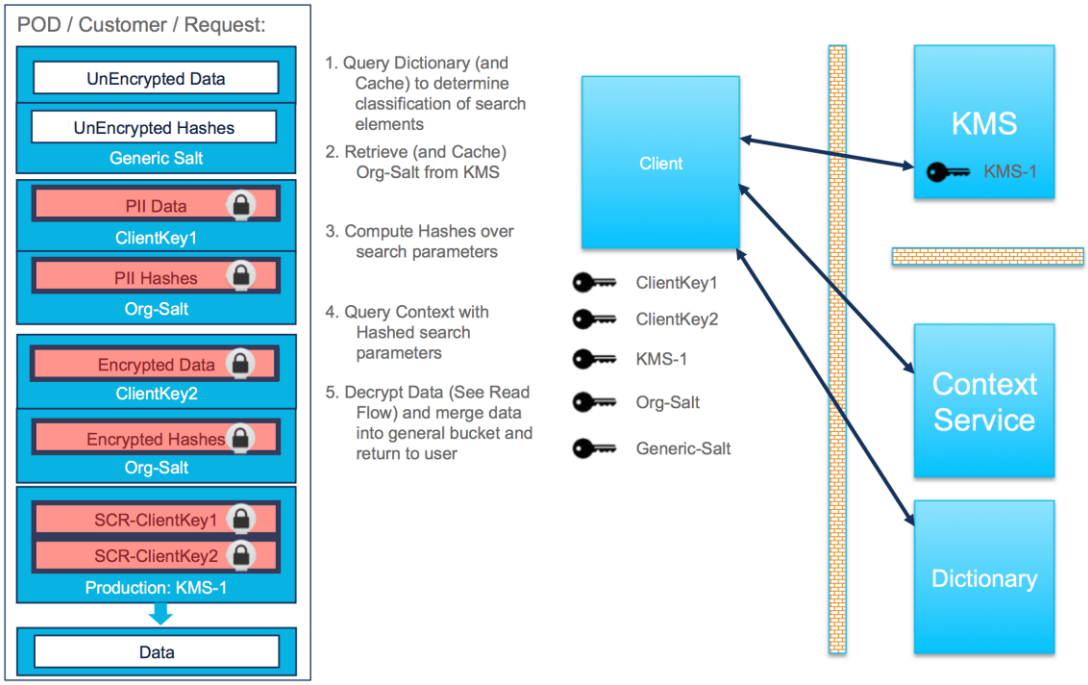

Partial Object Access

For partial object access, the user must be a member of the workgroup that can access the object and the workgroup for which access is being enabled. In the following diagram, the Production workgroup has access to PII and Encrypted data. Access is being enabled for just the Encrypted data for the Lab workgroup.

Figure 7: Enabling Partial Object Access to Additional Workgroups

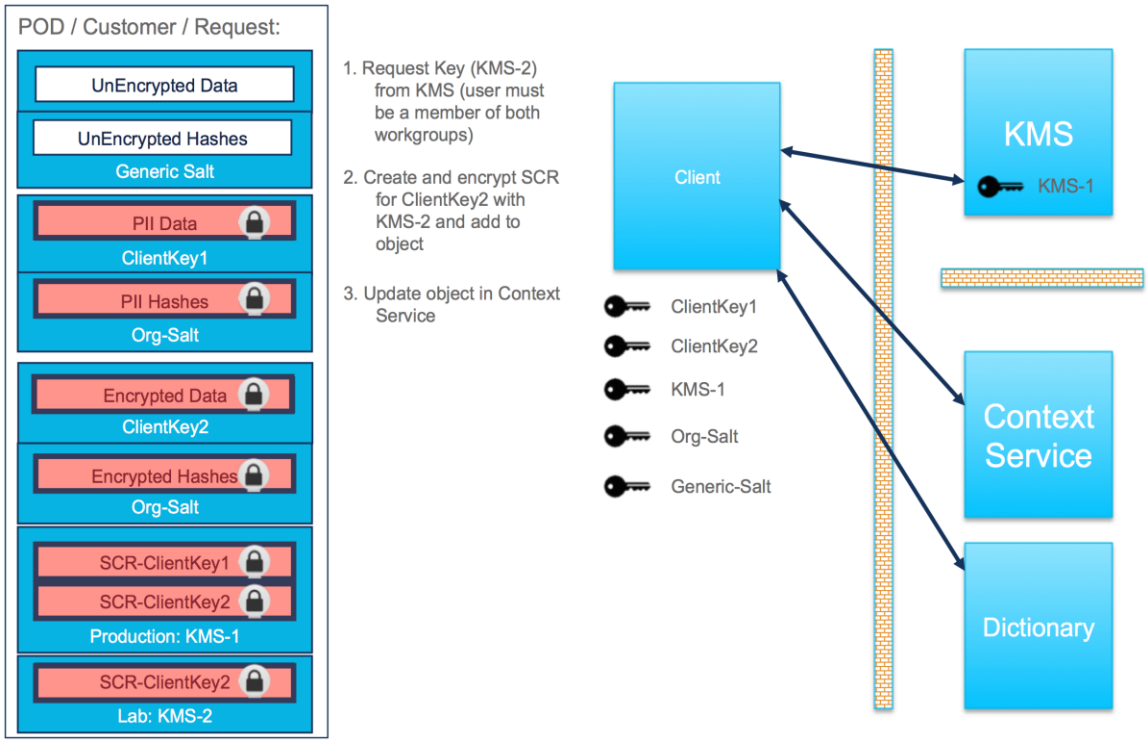

Access Control and Key Rotation

As previously described, access to data is controlled at the application layer (membership in workgroups) and re-enforced cryptographically. If a user is removed from a workgroup, access to the data is removed and the user is de-authorized from the corresponding KMS source.

To ensure that the user does not have access cryptographically to any new data, the KMS keys associated with the workgroups are rotated frequently. The current rotation schedule is 24 hours but is subject to change based on an organization's requirements.

Rotating the keys maintains the cryptographical re-enforcement across any data created after the user is removed from the workgroup. Without this, a user could have cached KMS keys while authorized to the KMS resource.

Summary

Context Service is a secure cloud-based storage service that provides a repository for customer journey data. It enables Cisco Contact Center customers to deliver a seamless omnichannel experience with an out-of-the-box integration from Cisco Customer Collaboration products.

Through the use of Cisco Common Identity and Cisco Spark KMS, Context Service provides true end-to-end encryption, allowing customer organizations to own their own security. It goes above and beyond the cloud industry standards of encryption of data in transport and at rest, limiting attack vectors beyond other industry solutions.

References

Cisco Spark Security and Privacy – http://www.cisco.com/c/dam/en/us/solutions/collateral/collaboration/cloud-collaboration/cisco-spark-security-white-paper.pdf