CMK Encryption on AWS (Preview)

Date published: 2021-08-05
Date modified: 2022-08-10

Legal Notice

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided ‘as is’ without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Contents

What is a Customer Managed Key (CMK)

In AWS, the Key Management System (KMS) provides keys for encrypting data and file systems. Typically, KMS is used to generate and manage keys, but as an alternative, you are able to generate a key yourself and use it to encrypt data and file systems in AWS. In this case, you are responsible for generating, maintaining, and rotating keys. This is the Customer Managed Key (CMK). You have full control over CMKs, including establishing and maintaining the key policies, IAM policies, grants, enabling and disabling them, rotating their cryptographic material, adding tags, creating aliases that refer to the CMK, and scheduling the CMKs for deletion.

CMKs are used to encrypt the following:

- EFS
- EBS volumes attached to nodes
- K8s etcd secrets
- S3 log buckets (that are used with CDE)

KMS keys are not used by CML to encrypt S3 buckets, which CDE uses and Liftie supports.

The key is automatically performed by Environment Services. For more information, see: CMK Encryption on AWS.

Limitations

- Only automatic key rotation is supported.
- CML only supports symmetric keys for CMK encryption

Implementing CMK on AWS

Follow the steps below to begin using CMK.

1. Update policies

To use CMKs in CML, ensure that the following three permission blocks are added to the policy section, in addition to the default policies:

```
{
   "Sid":"Allow Autoscaling service-linked role for attachment of
persistent resources",
   "Effect":"Allow",
   "Principal":{

"AWS":"arn:aws:iam::[YOUR-ACCOUNT-ID]:role/aws-service-role/autoscali
ng.amazonaws.com/AWSServiceRoleForAutoScaling"
   },
   "Action":"kms:CreateGrant",
   "Resource":"*",
   "Condition":{
      "Bool":{
         "kms:GrantIsForAWSResource":"true"
      }
   }
},
{
   "Sid":"Allow Autoscaling service-linked role use of the CMK",
   "Effect":"Allow",
   "Principal":{

"AWS":"arn:aws:iam::[YOUR-ACCOUNT-ID]:role/aws-service-role/autoscali
ng.amazonaws.com/AWSServiceRoleForAutoScaling"
   },
   "Action":[
      "kms:Encrypt",
      "kms:Decrypt",
      "kms:ReEncrypt*",
      "kms:GenerateDataKey*",
      "kms:DescribeKey"
   ],
   "Resource":"*"
}{
   "Sid":"Allow EKS access to EBS.",
   "Effect":"Allow",
   "Principal":{
      "AWS":"*"
   },
   "Action":[
      "kms:CreateGrant",
      "kms:Encrypt",
      "kms:Decrypt",
      "kms:ReEncrypt*",
```

```
            "kms:GenerateDataKey*",
            "kms:DescribeKey"
        ],
        "Resource":"*",
        "Condition":{
            "StringEquals":{
                "kms:CallerAccount":"[YOUR-ACCOUNT-ID]",
                "kms:viaService":"ec2.[YOUR-ACCOUNT-REGION].amazonaws.com"
            }
        }
}
```

2. Update the cross account role attached to the corresponding environment.

Overview of the setup steps

Configuring your environment to use a CMK involves the following steps:

1. Ensure that your provisioning credential has the minimum access permissions.
2. Ensure that your existing encryption key fulfills the required criteria or create a new encryption key according to the instructions provided in this document.
3. When creating an environment, specify the encryption key that should be used for encrypting the environment, including the Data Hubs running in it.

Note: Once your environment is running, if you would like to use a different key for encrypting a specific Data Hub, you can configure it as described in Encryption for Data Hub's EBS volumes on AWS.

Permissions for using encryption

If you are planning to use encryption, ensure that the cross-account IAM role used for the provisioning credential includes the following permissions:

EC2 permissions

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Effect": "Allow",
    "Action": [
```

```
      "ec2:CopyImage",
      "ec2:CreateSnapshot",
      "ec2:DeleteSnapshot",
      "ec2:DescribeSnapshots",
      "ec2:CreateVolume",
      "ec2:DeleteVolume",
      "ec2:DescribeVolumes",
      "ec2:DeregisterImage",
    ],
    "Resource": "*"
  }
}
```

KMS permissions

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Effect": "Allow",
    "Action": [
      "kms:DescribeKey",
      "kms:ListKeys",
      "kms:ListAliases"
    ],
    "Resource": "*"
  }
}
```

Encryption key requirements

If planning to use encryption, ensure that your encryption key can be used or create a new encryption key.

Ensuring that an existing encryption key can be used

If you already have an existing encryption key, make sure that the key fulfills the following requirements.

If you have an existing encryption key that you would like to use with Data Hub, make sure that:

- The following are attached as key user:
    - The *AWSServiceRoleForAutoScaling* built-in role.
    - Your IAM role or IAM user used for the cloud credential.
- To check that these are attached, in the AWS Management Console, navigate to the KMS console > Customer managed keys, select your encryption key, and scroll to Key Users.
- The encryption key is located in the same region where you would like to create clusters with encrypted volumes.

Create a new encryption key on AWS

If you don't have an existing encryption key, use the following instructions to create one.

1. In the AWS Management Console, navigate to KMS console.
2. Select Customer managed keys.
3. From the Region dropdown, select the region in which you would like to create and use the encryption key.
4. Click Create key.
5. In Step 1: Configure Key:
    1. Under Key type, choose Symmetric.
    2. Expand Advanced Options and under Key Material Origin, select “KMS” or “External”.
6. In Step 2: Create Alias and Description:
    1. Enter an Alias for your key.
    2. Defining Tags is optional.
7. In Step 3: Define Key Administrative Permissions, select the following:
    1. Choose your own IAM user / role used for logging into the AWS Management Console. Do not set *AWSServiceRoleForAutoScaling* or the cross-account IAM role as the key admin.
8. In Step 4: Define Key Usage Permissions:
    1. Select the *AWSServiceRoleForAutoScaling* built-in role.
    2. Select the cross-account IAM role.
9. In Step 5: Review and edit key policy, you may optionally tweak the key policy as desired, or simply leave it as generated by AWS.
10. Navigate to the last page of the wizard and then click Finish to create an encryption key.

Create an environment with a CMK

You can register your environment as described in Register an AWS environment from CDP UI, just make sure to specify the CMK that should be used to encrypt data, as described in the below steps.

Steps - CDP UI

1. Log in to the CDP web interface.
2. Navigate to the **Management Console > Environments**, and click **Register environment.**
3. Provide an **Environment Name.**
4. Select a **provisioning credential.**
5. Click **Next.**
6. Provide a **Data Lake Name.**
7. In the **Data Access and Audit** section, provide your data storage location and IAM resources created for minimal setup for cloud storage.
8. Click **Next.**
9. Select your region.
10. Under **Customer-Managed Keys**, click **Enable Customer-Managed Keys.**
11. In the same section, select the CMK:

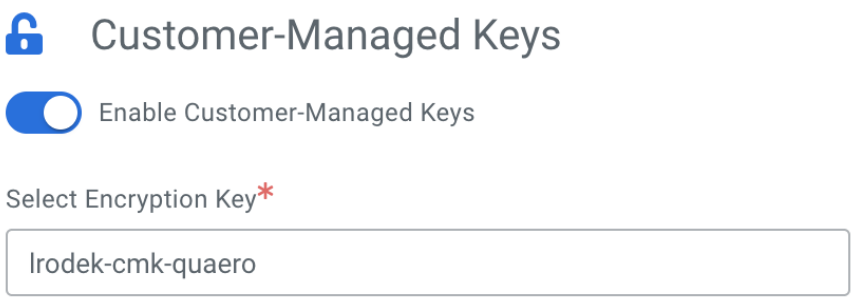

12. Select network, security groups, and provide an SSH key. If required, add tags.
13. Click **Next.**
14. In the **Logs** section, provide your logs storage location and managed identities created for minimal setup for cloud storage.
15. Click **Register Environment.**

Steps - CDP CLI

You can use your usual CDP CLI command to create an environment with a CMK, just add the **`--encryption-key-arn`** parameter and provide the encryption key created earlier. The easiest way to obtain the correct CLI template for creating an environment is by obtaining it from CDP wen UI as described in Obtain CLI commands from the register environment wizard.

For example:

```
cdp environments create-aws-environment \
--environment-name <ENVIRONMENT-NAME> \
--credential-name <EXISTING_CREDENTIAL> \
--region "<REGION>" \
--security-access cidr=<CIDR> \
--authentication publicKeyId="<SSH_KEY>" \
--log-storage storageLocationBase=<BUCKET_URL>,instanceProfile=<IDBROKER_IP>
\
--vpc-id <VPC_ID> \
--subnet-ids <SUBNETS \
--encryption-key-arn <ENCRYPTION_KEY_ARN>
```

The ARN of the encryption key created earlier should be passed in the parameter **--encryption-key-arn**

If the customer-managed encryption key ARN is not passed, then the AWS region-specific default encryption key is used for encrypting EBS volumes and RDS instances.

Set a CMK for an existing environment

You can set a CMK for an existing environment. The CMK is only used for encrypting disks of Data Hubs created after the CMK was added.

Note: The CMK added to an existing environment is only used for encrypting disks of Data Hubs created after the CMK was added. FreeIPA disks are not encrypted with the CMK. Data Lake and the external RDS are not encrypted except if they are created after the CMK was added.

Steps - CDP CLI

You can add an encryption key for an existing environment that does not yet have encryption enabled:

```
cdp environments update-aws-disk-encryption-parameters \
  --environment <ENVIRONMENT_NAME> \
  --encryption-key-arn <CMK_ARN> \
```

References

Instructions for using AWS IAM restricted Role and Policy for Compute(Liftie) & CML	https://docs.google.com/document/d/1zGIcYVF4Y8b8jsA2kvss5Qhb8WkS62peFpClBAovgM8
Customer master key - Concepts	https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html#master_keys
Rotating customer master keys	https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html
Creating CMK	https://docs.aws.amazon.com/kms/latest/developerguide/create-keys.html
How to choose your CMK configuration	https://docs.aws.amazon.com/kms/latest/developerguide/symm-asymm-choose.html