



Getting Started with Baffle Data Protection Services (AWS AMI)

Release 1.6.0.56

August 3, 2021

Overview

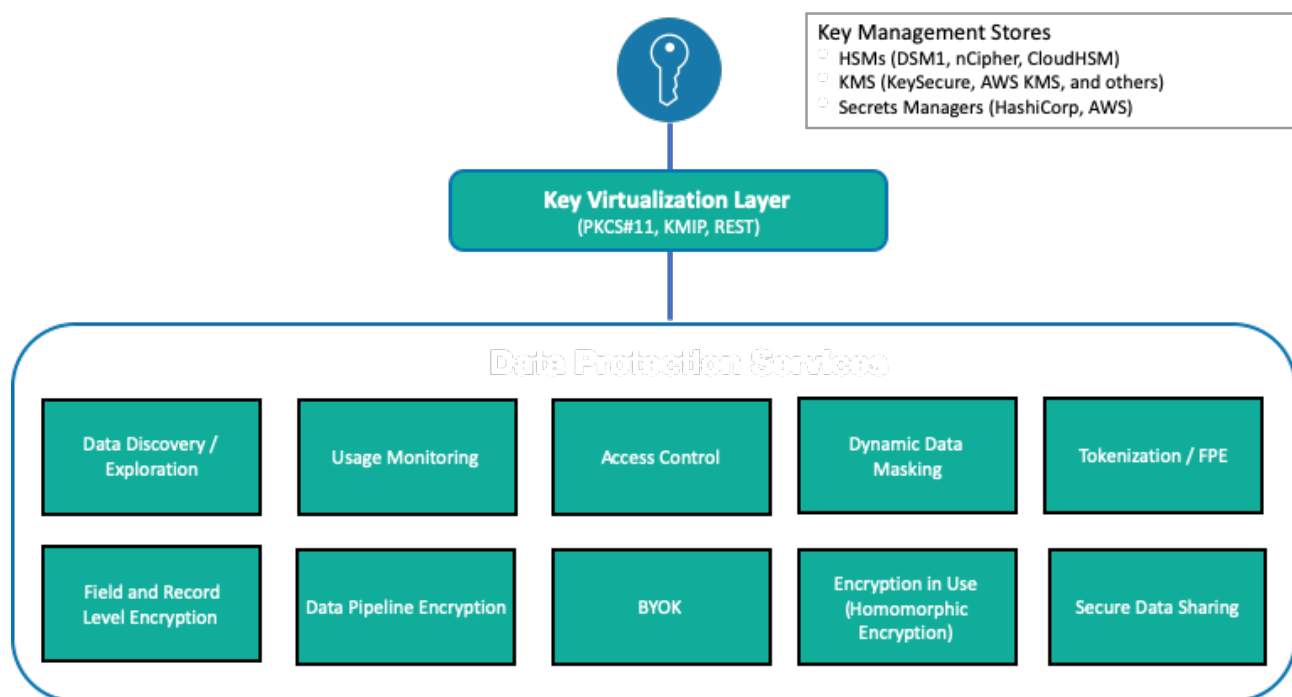
This guide gets you started using Baffle Data Protection Services in AWS. It describes Baffle Manager and Baffle Shield system requirements and architecture, followed by a configuration walkthrough to set up Baffle's column level encryption. The configuration steps are divided into five main tasks:

1. [Configure Baffle Manager](#) — configure the Baffle administrative console
2. [Connect to your Keystore](#) — configure the source for encryption keys
3. [Connect to your data store](#) — configure a connection to a database
4. [Configure a Baffle Shield](#) — configure the Baffle encryption engine
5. [Define a data protection policy for encryption](#) — configure a data protection policy

Background Information

Baffle Data Protection Services provide a range of data encryption, tokenization and de-identification methods to protect data in data stores and cloud storage environments. Common methods that Baffle employs include column or field level encryption, tokenization, format preserving encryption (FPE), dynamic data masking, and record level encryption.

Baffle integrates with key management stores via a key virtualization layer. It also provides for a local key store so you can use your own keys for data protection in the cloud.



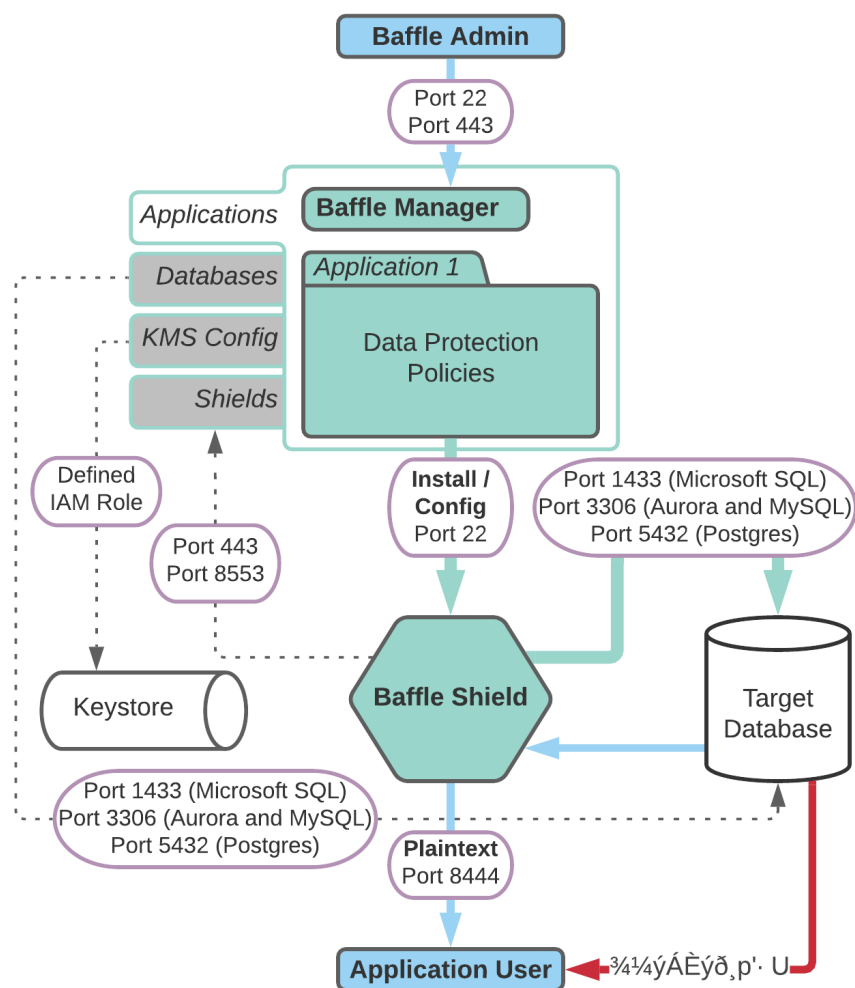
Pre-Requisites and Minimum System Requirements

Whether you use Baffle Professional Services to perform your deployment testing, or your organization does so independently as part of planning, ensure that your test environment meets the following minimum system requirements.

Baffle Component	Operating System	vCPU	Memory	Initial Space	Java
Baffle Manager	CentOS 7	2	8 GB	64 GB	OpenJDK Java 1.8
Baffle Shield	RHEL 7 or CentOS 7 equivalent	4	8 GB	64 GB ¹	OpenJDK Java 1.8
Database Platform	AWS RDS, Azure SQL and other supported database platforms ¹	16	256 GB	512 GB	OpenJDK Java 1.8
Prerequisite Information for Data Encryption					
Data Schema	<ul style="list-style-type: none"> Number of columns to be encrypted Data types and column field names Number of rows in table(s) Database size; Indexing, if any 				
Application	<ul style="list-style-type: none"> Identify the application and associated data for testing (for example, Microsoft SQL Server 2014 or later) Set aside a copy of the application and data to expedite troubleshooting and diagnostics. Provide test data that is encoded using UTF-8 character set. 				
Key Storage	<ul style="list-style-type: none"> Provide a supported key storage solution (see Key Management Support in the Baffle support center) Provide associated encryption keys Host in AWS and make available to Baffle infrastructure 				

¹ Additional supported database platforms are listed in the [Baffle support center](#).

Baffle Architecture and Communication



Port Requirements

Baffle Manager enables encryption policies and configurations by communicating with the Baffle Shield and your databases. Baffle Manager constructs a privacy schema that maps key IDs to data columns, thus enabling encryption in a simplified manner.

The following table lists the ports that must allow connections in order for Baffle Manager to communicate.

Host	Port Required	Direction	Purpose
Baffle Manager	22	Inbound	Console access for admin
Baffle Manager	443	Inbound	Web interface access for admin
Baffle Manager	8553	Inbound	Baffle Shield client access
Baffle Manager	22	Outbound	Baffle Shield configuration
Baffle Manager	1433	Outbound	Database schema mapping
Baffle Manager	5696	Outbound	(Optional) KeySecure access
Baffle Shield	22	Inbound	Console and Baffle Manager access
Baffle Shield	8444	Inbound	Application communication
Baffle Shield	1433	Outbound	Database access ¹
Baffle Shield	3306	Outbound	Database access ²
Baffle Shield	5432	Outbound	Database access ³
Baffle Shield	5696	Outbound	KeySecure access
Baffle Shield	8553	Outbound	Baffle Manager communications
Database Server ¹	1433	Inbound	Baffle Manager and Baffle Shield access

Database Server ²	3306	Inbound	Baffle Manager and Baffle Shield access
Database Server ³	5432	Inbound	Baffle Manager and Baffle Shield access
KeySecure	5696	Inbound	(Optional) Baffle Manager and Baffle Shield key config and retrieval

- ¹ For Microsoft SQL Server default port communications
- ² For MySQL, MariaDB or Aurora server default port communications
- ³ For PostgreSQL server default port communications

Configuration Prerequisites

Before you begin configuring Baffle Manager and Baffle Shield, verify that you have met the following requirements:

- AWS account with admin privileges
- SSH client
- Private key pair
- Database privileges for encryption and migration, see [Appendix A](#) for details.

Configuration Walkthrough (AWS)

To configure Baffle for your AWS environment, you will perform the following tasks. Click a link to jump to a task:

1. [Launch and configure the Baffle Manager AMI from AWS Marketplace](#)
2. [Connect to a Keystore](#)
3. [Connect to a Data Store](#)
4. [Launch and Configure a Baffle Shield AMI](#)
5. [Define a Data Protection Policy and Encrypt Your Data](#)

Task 1. Launch and configure the Baffle Manager AMI from AWS Marketplace

This task consists of two procedures. First, you will launch Baffle Manager and then configure Baffle Manager. **NOTE:** Baffle Manager AMI supports Docker containers.

To launch Baffle Manager, do the following:

1. Search for Baffle in the AWS Marketplace or click the following link to begin setup – [Baffle Data Protection Services](#), once on the page click **Subscribe**, then **Continue with Configuration**.
2. Make the following selections to **Configure this Software**:
 - a. Delivery Method – 64-bit (x86) Amazon Machine Image (AMI)
 - b. Software Version – Baffle Manager Release (latest version is displayed by default)
 - c. Region – Select the region.
3. Click **Continue Launch**, then under **Choose Action** select **Launch through EC2** from the drop-down list and click **Launch**.
4. Select the desired **Instance Type** from the list and click **Next: Configure Instance Details**.
5. On the Configure Instance Details page, accept the default settings with the exception of specifying the following:
 - a. Create a new security group on the VPC based on 'seller settings'. This configuration opens the necessary ports for Baffle Manager. Set the range of IP addresses that will be permitted access.
 - b. Ensure you have saved the selected key pair to access the Baffle Manager.
6. Click **Next: Add Storage**, enter the desired root storage **Size (GiB)** and click **Next: Add Tags**.
7. Click **Add Tag** and enter a **Key** and **Value**, then **Add another tag** with a Key and Value. Name and Owner are two commonly used tags.
8. Click **Next: Configure Security Group**. On the Configure Security Group page, accept **Assign a new security group**, along with the Security group name that provides the recommended ports and connection settings Baffle Manager requires for its data protection services.

IMPORTANT! You must add your own inbound security group with your IP address here so you will be able to connect to BM in the web browser.

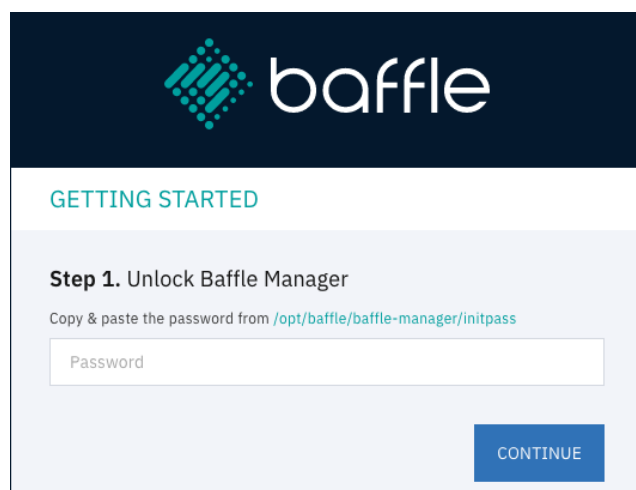
9. Click **Review and Launch** to review the instance configurations. Verify that the **PRIVATE_IP** address for Baffle Manager is correct in the `.env` file (located in the BM-Docker-Deploy directory), then click **Launch**.
10. You are prompted to **Select an existing key pair** or **Create a new key pair**. Once done, click the **I acknowledge** check box followed by **Launch Instances**.
11. Click **View Instances** to go to the EC2 dashboard. Enter one of the specified tags to search for the instance.

To configure Baffle Manager, do the following:

1. Once the instance is running, connect to it with a web browser via HTTPS. Use the public IP address of the instance, prefaced with `https://` for example, **`https://192.168.1.1`**.

NOTE: If you are unable to connect to the instance via HTTPS, check your security group inbound rules. Also ensure that your instance has finished initializing.

2. Since the instance is bootstrapped with a self-signed certificate, you will receive an invalid CA warning. Select the browser option to “proceed”. You will have the opportunity to upload and use your organization’s certificate later in this section. The following window appears.



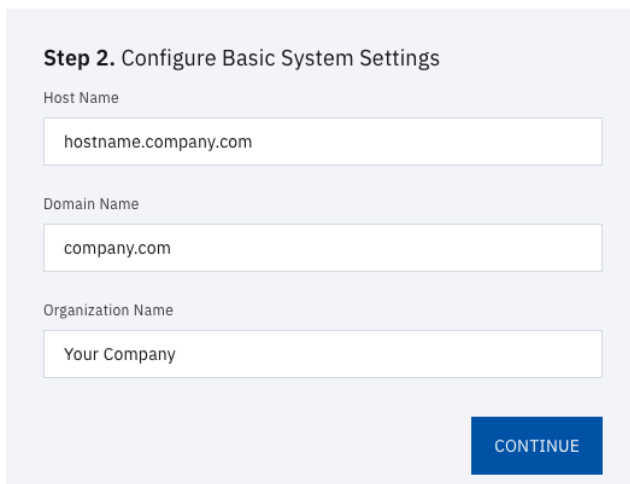
This dialog indicates that the Baffle Manager is in a locked state.

3. To unlock the Baffle Manager, access the system with SSH using “baffle” as the username for the SSH connection, followed by the public IP address (for example, `baffle@192.168.1.1`). You will also need the key pair file that you selected when you launched the instance. How you add the key pair depends on the shell client you are using (such as, SecureCRT).
4. You access the `initpass` file that unlocks Baffle Manager by connecting to the Baffle Manager instance with SSH, then using the following command to retrieve the unlock code.

```
sudo more /var/lib/docker/volumes/baffle_manager/_data/initpass
```


5. In the Unlock Baffle Manager dialog, paste the unlock code in the password field and click **CONTINUE**.
6. **Configure System Settings.** You are prompted for hostname and domain settings. All system users must have this domain name as part of this email going forward.

GETTING STARTED



Step 2. Configure Basic System Settings

Host Name

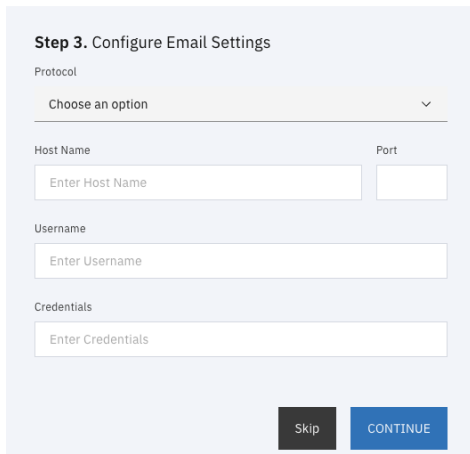
Domain Name

Organization Name

CONTINUE

7. **Configure Email Settings.** This allows Baffle Manager to send emails to provide notifications and for password resets. Enter the SMTP server to use, as well as the login credentials for the SMTP server.

GETTING STARTED



Step 3. Configure Email Settings

Protocol

Choose an option ▾

Host Name Port

Enter Host Name

Username

Enter Username

Credentials

Enter Credentials

Skip **CONTINUE**

8. **Create Admin Account.** The screen below prompts you to create the initial Baffle Manager administrator account. This account is used to configure the subsequent components such as the key management store, data store connections, and Baffle Shields.

GETTING STARTED

Step 4. Create Baffle Manager Admin User

Email Address

First Name

Last Name

Phone Number

Password

Confirm Password

At least 10 characters or longer. A mixture of both uppercase and lowercase letters. A mixture of letters and numbers.

CONTINUE

9. **Configure Credential Keystore.** This configuration screen establishes an encrypted credential store for any system access credential or access key that the Baffle Manager or Baffle Shield utilize. The default name is “baffle_credential_store” and cannot be changed.

Select LOCAL for Keystore type. For Secret Key, enter any random string which will be used to generate a random key to encrypt the Keystore Config Password. For Config Password, enter a secure password or passphrase to secure the actual keystore.

GETTING STARTED

Step 5. Configure Credential Keystore

Keystore Name

Keystore Type

Choose an option ▼

Baffle Secret Key

Config Password


Confirm Config Password


CONTINUE

10. **Install SSL Certificate.** This configuration step allows you to install an SSL certificate to secure access to the Baffle Manager web interface. Upload the certificate and key file for your organization or respective CA to enable SSL for the Baffle Manager console.

GETTING STARTED


Step 6. Configure HTTPs Certificate

Select Certificate File 

Select Key File 

Skip CONTINUE

11. This should complete the initial setup process and bring you to the login page.

 Baffle Manager setup was successful. Please login to use Baffle Manager.

Username

Password [Forgot Password?](#)

SIGN IN

12. Enter the credentials for the administrator account you created in Step 9 to login and continue the configuration process.

- ## Task 2. Connect to a Keystore

Example of a Local Keystore configuration:

ADD KEYSTORE ×

Keystore Name

Description

0 / 100

Keystore Type

LOCAL

▼

Baffle Secret Key

Cancel

Add Keystore

Example of an AWS KMS configuration:

ADD KEYSTORE ×

Description

0 / 100

Keystore Type

AWS_KMS

▼

AWS Access Key ID (Optional)

AWS Secret Access Key (Optional)

AWS Root ARN (Optional)

AWS Region

us-west-2

▼

DEK Storage Type

AWS S3 bucket

▼

App Namespace

Bucket Name

Cancel

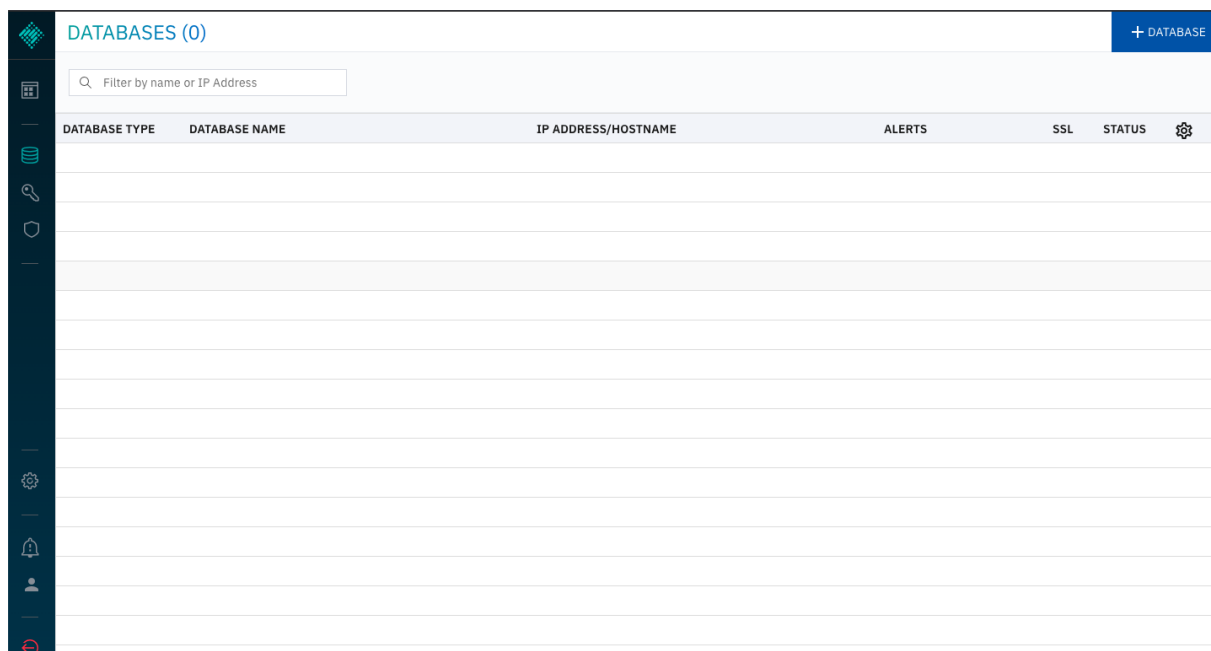
Add Keystore

Task 3. Connect to a Data Store

In this section, you will configure a connection to a database. This connection will allow Baffle Manager to enumerate fields or columns that can be selected as part of a data privacy policy, in order to enable column level encryption.

To connect to a data store, do the following:

1. **Display the list of configured databases.** Click on the database icon on the left navigation panel to display a list of configured databases.



2. **Enroll a database.** Click **+DATABASE** to add a Data Store. Enter a database name and description.
 - a. Specify the database type, and enter the hostname or IP (endpoint) for the database. For AWS RDS enter the endpoint URL.
 - b. Enter the port for the database. Default database ports are listed in the [Port Requirements](#) table.
 - c. Enter the database user credentials. **It is recommended that you create a new user on your database for use with Baffle.** See [Appendix A](#) for details.
NOTE: For information on how to allow users on your database with less privileges to access the encrypted data, see [Appendix B](#).
 - d. Select Use SSL to enable an SSL/TLS connection to the database.

The following example is of a Microsoft SQL Server configuration.

ADD DATABASE
X

Database Name
MS SQL Server

Database Description (Optional) 0 / 100
Short description

Database Type
RDS-SQL Server

Hostname/IP Address
hb-sqlserver-temp.c9mc! 1433

Database Username
baffleuser

Database Credential
.....

☐ Use SSL

Add file

Cancel Add Database

For Postgres Configurations:

You must enter the name of the database that you intend to connect in the field 'Postgres Database Name'. Otherwise, Baffle Manager will connect to a database with the same name as your database username. If no such database exists, then the connection attempt will fail.

The default database name is "postgres".

Database Type
Postgres

Hostname/IP Address

Port
5432

Database Username

Database Credential

Postgres Database Name
postgres

Task 4. Launch and Configure a Baffle Shield AMI

This section walks through the installation and configuration of a Baffle Shield. The Shield will be used to enforce a Data Protection Policy, encrypting the data in the databases that were configured in the previous section.

To launch Baffle Shield, do the following:

1. **Configure an AMI instance** to run the Baffle Shield.
 - a. In AWS, go to EC2 and launch a new AMI instance with a CentOS 7 operating system and appropriately sized for your environment.
 - b. Issue the following bootstrap commands in the Advanced Details section during the instance setup process.

```
#!/bin/bash
sudo su
yum install java-1.8.0-openjdk-devel -y
yum install mysql -y
yum install nano -y
yum install postgresql -y
yum install unzip -y
curl "https://awscli.amazonaws.com/awscli-exe-linux-x86_64.zip" -o "awscliv2.zip"
unzip awscliv2.zip
./aws/install
```

▼ Advanced Details

Metadata accessible ⓘ	Enabled
Metadata version ⓘ	V1 and V2 (token optional)
Metadata token response hop limit ⓘ	1
User data ⓘ	<input checked="" type="radio"/> As text <input type="radio"/> As file <input type="checkbox"/> Input is already base64 encoded

```
#!/bin/bash
sudo su
yum install java-1.8.0-openjdk-devel -y
yum install mysql -y
yum install nano -y
yum install postgresql -y
yum install unzip -y
curl "https://awscli.amazonaws.com/awscli-exe-linux-x86_64.zip" -o "awscliv2.zip"
unzip awscliv2.zip
./aws/install
```

- c. Select the same security groups you used for the Baffle Manager configuration. Ensure the security group for your Baffle Shield allows inbound connections from Baffle Manager (on port 22) and from your own IP address (on port 8444 by default).
 - d. Once you complete the setup process, allow the instance a few minutes to initialize.
2. **Connect the Baffle Shield to Baffle Manager.** Once the instance is running, return to your Baffle Manager admin interface. Click on the shield icon on the left navigation panel. A list of connected Baffle Shields appears. Click **+BAFFLE SHIELD** in the upper right corner.

2. Do the following:
 - a. Select “Automated Deployment” for Deployment Model.
 - b. Enter the Host Username “centos” to access the Baffle Shield EC2 Instance.
 - c. Enter the IP Address of the Baffle Shield you just launched. **NOTE:** If your Shield runs in the same VPC as your Baffle Manager instance, it is recommended that you use the Private IP address here.
 - d. Enter a port number that the Baffle Shield will use to listen for application connections. The default port is 8444.
 - e. Select “Use SSL” if the data store connection uses SSL.
 - f. Select “Use SSH Key” and upload the key that you selected when you set up the Shield instance.
 - g. Optionally, a username and password can be used to access the Baffle Shield.
3. Click **Add Baffle Shield** to complete the process. The new Shield is added to the list of configured Baffle Shields.

NOTE: If the Baffle Manager is unable to connect to the shield, verify that your Shield's security group permits inbound access from Baffle Manager.

[illegible]

ENROLL APPLICATION
×

Application Name

Application Description (Optional) 0 / 100

Baffle Shields

Datastore

Keystore

Workload Capture
☐ Off

Encryption Method

2. **To Enroll Application**, enter a name and description and do the following:
 - a. Choose the Baffle Shield from the drop down that was configured in the previous section.
 - b. Select the Data Store which you will encrypt.
 - c. Select the Keystore to be used as a source for data encryption keys.
 - d. Specify the operational mode for the Baffle Shield. Leave Workload Capture Off, unless profiling an application.
 - e. Specify Column Level for the Encryption Method.
 - f. Click **Enroll Application**.

The following is an example of enrolling an application and deploying a Data Protection Policy for a MySQL database.

ENROLL APPLICATION

Application Name

MySQL Application 01

Application Description (Optional)

65 / 100

My policy plan: encrypt first five columns of table 'superstore'.

Baffle Shields

1 x Choose an option

Keystore

localkeystore

Datasource

MySQL Database 01


Workload Capture

☐ Off

Encryption Method

Column Level

Upload Entity Schema



Cancel

Enroll Application

The Applications page displays the new application.

[illegible]

3. **Define the Data Protection Policy.** Select the Application you just configured, then click **ENCRYPT** in the side bar that appears showing information about the application.

MySQL Application 01
⚙️ ×

🔒 Encrypt

DETAILS
Added on: 2020-10-27 7:52:58
Created by: devops@baffle.io

DESCRIPTION
My policy plan: encrypt first five columns of table 'superstore'.

ENCRYPTION DETAILS
Enc Type: Column Level
Enc Mode: Classic
Key Rotation: 0
Database Name: MySQL Database 01
keystore: localkeystore

MIGRATION DETAILS
Migration Plan: Same Database
Batch Size: 2000
Failure Scope: SERVER

ADVANCED CONFIGURATION
Workload Capture
☐ Off
Filter Mode
☐ Off

IP FILTERING ✎
Permitted IP Addresses

Blocked IP Addresses

BAFFLE SHIELDS

The Policy Builder window appears for the configured data store

POLICY BUILDER: MySQL Application 01
×

0 COLUMN(S) SELECTED FOR ENCRYPTION

Filter by database name

Filter by table name

Filter by column name

DATABASE	COL SEL	TABLE	COL SEL	🔒 COLUMN N	DATA TYPE	ENC MODE	DATA F
test3							

Cancel Next

4. **Select the database and table to encrypt.**

POLICY BUILDER: MySQL Application 01

5 COLUMN(S) SELECTED FOR ENCRYPTION

category city country customerid customername

Filter by database name Filter by table name Filter by column name

DATABASE	COL SEL	TABLE	COL SEL	ENC MODE	DATA FORMAT	KEY ID
test3	5	superstore	5	AES-CTR-DET	OFF	2
				AES-CTR-DET	OFF	2
				AES-CTR-DET	OFF	2
				AES-CTR-DET	OFF	2
				AES-CTR-DET	OFF	2
				AES-CTR-DET	OFF	2
				AES-CTR-DET	OFF	2
				AES-CTR-DET	OFF	2
				AES-CTR-DET	OFF	2
				AES-CTR-DET	OFF	2

Cancel Next

5. **Select columns** for encryption and the respective encryption mode.
6. Optional: Specify Key IDs for use to encrypt specific columns. Scroll to the right on the column selector and add more keys by clicking (+). The default value for Key ID is 2. Available Key IDs will be displayed in the Key ID dropdown menu for each column.

POLICY BUILDER: MySQL Application 01

5 COLUMN(S) SELECTED FOR ENCRYPTION

category city country customerid customername

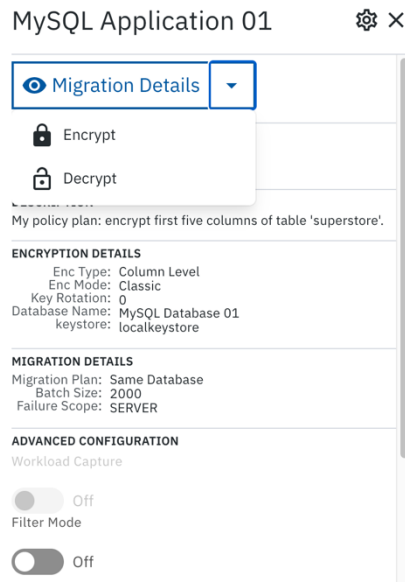
Filter by database name Filter by table name Filter by column name

DATABASE	COL SEL	TABLE	COL SEL	COLUMN N	DATA TYPE	ENC MODE	DATA F
test3	5	superstore	5	<input checked="" type="checkbox"/> category	varbinary(248)	AES-CTR-DET	OFF
				<input checked="" type="checkbox"/> city	varbinary(368)	AES-CTR-DET	OFF
				<input checked="" type="checkbox"/> country	varbinary(368)	AES-CTR-DET	OFF
				<input checked="" type="checkbox"/> customerid	varbinary(248)	AES-CTR-DET	OFF
				<input checked="" type="checkbox"/> customern	varbinary(568)	AES-CTR-DET	OFF
				<input type="checkbox"/> orderdate	varbinary(163)	AES-CTR-DET	OFF
				<input type="checkbox"/> orderid	varbinary(192)	AES-CTR-DET	OFF
				<input type="checkbox"/> productid	varbinary(192)	AES-CTR-DET	OFF
				<input type="checkbox"/> productnan	varbinary(592)	AES-CTR-DET	OFF
				<input type="checkbox"/> region	varbinary(192)	AES-CTR-DET	OFF

Cancel Next

- [illegible]

9. To Decrypt data, click on the application again, and select DECRYPT from the dropdown menu. The Policy Builder will re-open. Select the columns which you would like to decrypt and click NEXT to proceed. Only the columns that you have previously encrypted will be available to decrypt.



Summary

You have now completed configuration of the Baffle Manager, Baffle Shield and created a Data Protection Policy to protect your data.

To confirm your data is encrypted, access the database normally with your SQL client. You should find the columns you selected are now encrypted.

To view the columns in the clear, use your SQL client to connect to the Baffle Shield. Connect using the public IP address of the Shield, port 8444, and the credentials for the database user you submitted in [section 3, step 2b](#). Access the encrypted tables, and you should find the columns are visible.

Appendix

Task A: Database Privileges for encryption and migration

In order to carry out encryption and migration, Baffle Shield requires certain user permissions on the database. It is recommended that you create a new user on your database for Baffle Shield to use, rather than assign your database administrator.

Use your SQL client to issue the following grants. Enter the credentials of this new user in [section 3, step 2b](#), so that Baffle Shield has full privileges to encrypt and decrypt the data you select.

1. To create a new user:

- a. **create user** '<baffle user>'@'%' ;
- b. **set password for** '<baffle user>' =
password('<password>') ;

2. To grant the requisite permissions:

- a. **GRANT USAGE ON *.* TO** '<baffle user>'@'%' ;
- b. **GRANT ALL PRIVILEGES ON** shadow_information_schema.* **TO**
'<baffle user>'@'%' ;
- c. **GRANT ALL PRIVILEGES ON** <target database>.* **TO** '<baffle user>'@'%' **WITH GRANT OPTION;**

Repeat step c for each database you wish to encrypt. When completed, Baffle Shield has the necessary permissions in order to carry out encryption and migration. Use the credentials of the user specified here.

Task B: Minimum Required Database Privileges

These are the minimum required grants for users on your database who need the least access privileges. Use your SQL client to issue the following commands with your admin user. These grants permit the restricted-access user to obtain only the data you specify.

For MySQL and Aurora databases:

1. Issue the following commands.
 - a. **GRANT USAGE ON *.* TO '<username>'@'%' ;**
 - b. **GRANT ALL PRIVILEGES ON shadow_information_schema.* TO '<username>'@'%' ;**
 - c. **GRANT SELECT ON <target database>.<target table> TO '<username>'@'%' ;**
 - d. Repeat step c for each table you wish to make accessible to the user. When completed, you may connect to the Baffle Shield proxy with this user.
 - e. To confirm user privileges, use: **show grants ;**
2. OPTIONAL: Some databases may require additional information from the user. Take the hash of the user's password with the following:
 - a. **SELECT PASSWORD ('<user password>') ;**
Insert the hash back into the expressions:
 - b. **GRANT USAGE ON *.* TO '<username>'@'%' IDENTIFIED BY PASSWORD '<password hash>' ;**
 - c. **GRANT ALL PRIVILEGES ON shadow_information_schema.* TO '<username>'@'%' IDENTIFIED BY PASSWORD '<password hash>' ;**
 - d. **GRANT SELECT ON <target database>.<target table> TO '<username>'@'%' IDENTIFIED BY PASSWORD '<password hash>' ;**
 - e. Repeat step d for each table you have selected.