


Container

AWS - INSTALLATION & CONFIGURATION GUIDE



Contents

Introduction.....	4
Requirements.....	6
Step A - ECS – Provision Cludockit Container (hosted on ECS/Fargate)	7
STEP 0 – Create a S3 Bucket and upload the license	7
STEP 1 – Create an AWS Secret in SECRET MANAGER to store information to connect to Cludockit Container Repository	7
STEP 2 - Creating Policy to access the Secret.....	7
STEP 3 - Creating Role assumed by ECS tasks to access the Secret.....	8
STEP 4 - Creating VPC and Subnet + Open required ports	8
STEP 5 - Creating Security Group.....	8
STEP 6 - Creating the Task Definition File	8
Parameters.....	8
Environment variables.....	8
STEP 7 - Creating the Task Definition.....	10
STEP 8 - Creating the cluster and register Task Definition.	10
STEP 9 – Validate Service is Up & Running.....	10
Step B (Optional) – Configure Cludockit Web UI.....	11
Optional – Store these credentials in AWS Secrets Manager.....	15
Step C (Optional) – Configure Cludockit Container to support Scheduling.....	16
Start Cludockit Scheduler Container.....	16
Set Settings in the settings file	16
Step D (Optional) – Configure Cludockit Container to support the creation of Compliance Rules, Tailored Diagrams and Settings	17
Create (or re-use) an Azure Cosmos DB.....	17
Configure Cludockit Container to use the Azure Comos DB.....	19
Step E – Understand Cludockit API Container	20
Step F – Test your license.....	21
Activate and setup components for your license.....	21
Step G – Validate that you can authenticate to the environment that you want to scan.....	22

Step H – Test the document generation	24
Step I – Manage your document generation.....	25
/ListDocumentGeneration.....	25
/StopDocumentGeneration	25
Annex – Deploy multiple instances of Clouddokit Container	26
Step 1 – Create / Configure Azure Key Vault	27
Step 2 – Configure Azure Redis Cache	28
Step 3 – Define the Environment Variables required to run the Clouddokit Container	28
Annex - AWS Container use-cases.....	32
Introduction.....	32
AWS ECS Container	32
A – Scanning multiple AWS accounts using Cross-Account role.....	32
B – Using the container with the ECS Task Role instead of IAM user keys	37
C – Running a scheduled scan in AWS using ECS container	39
Container details.....	42
Environment Variables.....	42
Logging.....	43
Annex – Troubleshooting	44

Introduction

The purpose of this document is to provide the detailed steps to run and configure Clouddokit Docker container image.

There are two types of images that you should run:


- **cdk-web-linux** that contain the Clouddokit API/Web interface. This is mandatory to run this container.
- **cdk-scheduler-linux** that contain the Clouddokit Scheduling features. This is an optional container you do not need to install if you do not want to use schedules.

The cdk-web image contains the Clouddokit API that you can call from your CI/CD processes or any other process / scenario which fits your business needs.

In addition to the API, we have integrated the complete Clouddokit Web UI in the image so that you can get all the features that you are accustomed to.

Clouddokit Docker container images provide you a way to run Clouddokit into your own isolated Cloud environment and gives you the exact same features as Clouddokit Website and Clouddokit Desktop.

Here is the high-level overview of the solution:



The following hosting environments are currently supported:

- Web App for Containers on Azure - Recommended
- ECS (Elastic Container Services) Fargate on AWS
- ACI (Azure Container Instance) on Azure
- GKE (Google Kubernetes Engine) on GCP

A few important things to note:

- These configurations are for the hosting of the container, not for the environment that you scan which means that you can scan a GCP project using the Clouddokit Container API even if the container runs on Azure.

- Depending on the hosting option that you choose, there could be some limitations. Those limitations are related to the hosting option and not the Cludockit Container itself. As an example, ACI currently does not yet support private networking (virtual networks) for Windows Based Container.
- The current document does not detail networking configuration like isolation/https setup as this is highly depending on your internal setup.
- Container is currently designed to have one node running which should be more than enough to generate all your documents you need.
- For production environment, we recommend 4vCPU + 8 Gb RAM
- Cludockit Web UI only supports Azure AD as SSO authentication. If you do not set it up, you will only be able to access the API portion.

The following sections contains the different steps to deploy the Cludockit Docker container image on the AWS Platform.

Here is an overview of the different steps you must do to deploy Cludockit Container:

Step A - Deploy Cludockit Container (cdk-web)

- This Step is subdivided in 8 Different steps
- Those steps have been scripted to ease deployment process so we strongly advice to use the scripted approach

Step B (Optional) - Activate Cludockit Container UI

- Create an Azure AD Application

Step C (Optional) - Activate the Scheduling feature (cdk-scheduler)

- Set the appropriate settings to activate scheduling

Step D ... - Do some tests

- Test the license validity
- Start some documentation

Requirements

To install Cludockit Container in your environment, you will need:

- A S3 Bucket
- A Secret Manager to store the secret to pull the image from Cludockit Repo.
- A Role / Policy to access the Secret.
- A VPC/Subnet to host the ECS Cluster.
- A ECS Cluster (Fargate)
 - Note that this is the recommended approach as this is the one we have scripted but it does not mean that we do not support other types of deployment like AKS or other as Cludockit Container is a standard plain container.
- (Optional) An Azure Active Directory Application if you want to activate Cludockit Container Web UI

Important note

We highly recommend that you use the script (based on AWS CLI) provided by Cludockit Team to provision/test Cludockit Container.

Step A - ECS – Provision Cludockit Container (hosted on ECS/Fargate)

STEP 0 – Create a S3 Bucket and upload the license

To run Cludockit Container, you need to have a S3 Bucket that will be used to store information (license file, settings...). As the license file is linked to this S3 bucket, you need to choose a container Name and send that name to Cludockit Support Team (support@cludockit.com) so that they generate a license file.

Once you receive your license file, you can upload the license file into a file named ***cludockitinternal/license.json***.

Please note that the Json license file is tied to the S3 Bucket name so you need to create/use an S3 Bucket with a name that matches the name provided to Cludockit Support Team.

Important note

Ensure that the Storage Account exists in your environment before sending it to the Cludockit Support Team.

Code extract

Refer to STEP 0 in the provided script for **ECS Fargate**.

STEP 1 – Create an AWS Secret in SECRET MANAGER to store information to connect to Cludockit Container Repository

In this step, you need to create a Secret that will contain the information to connect to Cludockit Container Registry so that your ECS can pull the image.

Secret should contain *username* and *password* properties, see code for reference.

Code extract

Refer to STEP 1 in the provided script for **ECS Fargate**.

STEP 2 - Creating Policy to access the Secret

Once you have created the secret, you need to create a policy that will allow access to the secret.

The policy needs to provide **Allow Effect** to the action **secretsmanager:GetSecretValue** to the secret previously created.

Code extract

Refer to STEP 2 in the provided script for **ECS Fargate**.

STEP 3 - Creating Role assumed by ECS tasks to access the Secret

Once the policy has been created, we need to create a Role attached to that policy so that ECS assumes the role to be able to pull the remote image.

The role needs to have the **Allow Effect** and needs to have the action **sts:AssumeRole** for the service *ecs-tasks.amazonaws.com*

Code extract

Refer to STEP 3 in the provided script for **ECS Fargate**.

STEP 4 - Creating VPC and Subnet + Open required ports

In this step, we need to configure a VPC, a Subnet, and a Route Table so that the subnet is reachable and so that the subnet has connectivity to the Internet to pull the image. Please note that you can use an existing VPC if it has the appropriate requirements.

Code extract

Refer to STEP 4 in the provided script for **ECS Fargate**.

STEP 5 - Creating Security Group

To Secure the deployment, we create a security group with only port 8080 open (please use https/443 for Production usage)

Code extract

Refer to STEP 5 in the provided script for **ECS Fargate**.

STEP 6 - Creating the Task Definition File

Once networking configuration is done, we need to create a task definition file that contains the following information.

Parameters

Name	Value
family	Clouddockit
executionRoleArn	arn:aws:iam::<acctnumber>:role/ecsTaskExecutionRole
networkMode	Awsvpc
requiresCompatibilities	Fargate
ContainerDefinitions	Should refer to the image and the secret to access it

Environment variables

Name	Value
AppInsightsConnectionString (optional)	An Azure Application Insights Connection String for advanced logging

Name	Value
DockerStorageCloudProvider	Specify if your Storage Account is stored in Azure, AWS or GCP. Possible values are: • Azure • GCP • AWS (select this value)
DockerStorageAWSBucketName	Enter the name of the S3 Bucket
DockerStorageAWSAccessKeyId (option 1)	Enter the Access Key Id for Full control of the AWS S3 Bucket
DockerStorageAWSSecretAccessKey (option 1)	Enter the Secret Access Key for Full control of the AWS S3 Bucket
DockerStorageUseAwsRole (option 2)	True (permissions are provided by the ECS that hosts the container)
DockerStorageUseAwsGov (optional)	Specify "True" if you are using Aws Gov Cloud or "False" otherwise. If the variable is not present, it's "False" by default
ExternalBaseUrl (optional)	Full public URL of the container when running behind a reverse proxy or load balancer (e.g. https://cloudockit.company.com). Required for OAuth redirect URIs to resolve correctly. When not set, the URL is derived from request headers only.
ExternalBaseUrlProxyHeader (optional)	Name of an HTTP header whose presence signals a proxied request (e.g. X-Forwarded-Proto). When set, the ExternalBaseUrl override applies only to requests carrying this header; other requests use the Host header.
TrustAllProxies (optional)	Set to True to trust all X-Forwarded-* headers without a URL override. Use only when the container is behind a trusted proxy and is not directly internet-exposed. Prefer ExternalBaseUrl.

Name	Value
ForwardProxyHopLimit (optional)	Number of proxy hops to trust for X-Forwarded-For. Default: 1. Increase when multiple load balancers sit in front of the container.

➔ **Note: parameters noted as option 1 and 2 are mutually exclusive**

Code extract

Refer to STEP 6 in the provided script for **ECS Fargate**.

STEP 7 - Creating the Task Definition

Then, from the Task Definition file, we just create a task definition:

Refer to STEP 7 in the provided script for **ECS Fargate**.

STEP 8 - Creating the cluster and register Task Definition.


Last step is to create a cluster and start the service:

Refer to STEP 2 in the provided script for **ECS Fargate**.

STEP 9 – Validate Service is Up & Running

From the AWS Console, navigate to the cluster, tasks and then click on the public IP to open it.

You should see Clouddokit Container running:



Step B (Optional) – Configure Clouddokit Web UI

Clouddokit Container supports a Web UI that allows users to authenticate by using Azure AD or Azure User Authentication.

This Web UI supports Azure Active Directory as a first step to authenticate users.

Once connected, you will be able to connect to Azure, AWS and GCP using Service Accounts (Azure AD App, GCP Service Credentials, AWS Access Keys).

To activate Azure AD Authentication, you need to follow these steps:

- Go to your **Azure Active Directory**
- Click on **App Registration** and then click **New Registration**
- Enter a **Name** (any name you want) and select **Single Tenant**
- Enter the following **redirect URIs** (reply url):
 - `https://<AppSvcName>.azurewebsites.net/LogIntoAzure/CatchCodeAzure`
 - `https://<AppSvcName>.azurewebsites.net/LogIntoCDKWithAAD/CatchCode`where *AppSvcName* is the name of your App Service

Register an application ...

* Name
The user-facing display name for this application (this can be changed later).

Clouddokit Web UI ✓

Supported account types
Who can use this application or access this API?

☒ Accounts in this organizational directory only (Single tenant)

☐ Accounts in any organizational directory (Any Azure AD directory - Multitenant)

☐ Accounts in any organizational directory (Any Azure AD directory - Multitenant) and personal Microsoft accounts (e.g. Skype, Xbox)

☐ Personal Microsoft accounts only

[Help me choose...](#)

Redirect URI (optional)
We'll return the authentication response to this URI after successfully authenticating the user. Providing this now is optional and it can be changed later, but a value is required for most authentication scenarios.


Web ▼ `https://contoso.azurewebsites.net/LogIntoAzure/CatchCodeAzure` ✓

Register an app you're working on here. Integrate gallery apps and other apps from outside your organization by adding from [Enterprise applications](#).

By proceeding, you agree to the [Microsoft Platform Policies](#)

Register

Note: the interface will not let you enter the 2nd URL before clicking on **Register** so you'll have to enter it after registration, in the Authentication page:



Then, go to **API Permissions**, click on **+Add a permission** and select:

- **Microsoft Graph**, then **Delegated permissions** and then select **User.Read**:
- **Azure Service Management**, then **Delegated permissions** and then select **user_impersonation**:

Request API permissions

[All APIs](#)

 Microsoft Graph
<https://graph.microsoft.com/> [Docs](#)

What type of permissions does your application require?

Delegated permissions
Your application needs to access the API as the signed-in user.

Application permissions
Your application runs as a background service or daemon without a signed-in user.

Select permissions [expand all](#)

 The "Admin consent required" column shows the default value for an organization. However, user consent can be customized per permission, user, or app. This column may not reflect the value in your organization, or in organizations where this app will be used. [Learn more](#)

Permission	Admin consent required
> IdentityRiskyUser	
User (1)	
☒ User.Read ⓘ Sign in and read user profile	No
☐ User.Read.All ⓘ Read all users' full profiles	Yes
☐ User.ReadBasic.All ⓘ Read all users' basic profiles	No
☐ User.ReadWrite ⓘ Read and write access to user profile	No

Add permissions

Discard

INSTALLATION & CONFIGURATION GUIDE

13

Click **Add permissions**. You should now see the following:

[Refresh](#) | [Got feedback?](#)

i The "Admin consent required" column shows the default value for an organization. However, user consent can be customized per permission, and user consent will be used. [Learn more](#)

Configured permissions

Applications are authorized to call APIs when they are granted permissions by users/admins as part of the consent process. The list of configured permissions should include all the permissions the application needs. [Learn more about permissions and consent](#)

[+ Add a permission](#) [✓ Grant admin consent for UMAknow Solutions DEV Inc](#)

API / Permissions name	Type	Description	Admin consent req...
▼ Azure Service Management (1)			
user_impersonation	Delegated	Access Azure Service Management as organization use...	No
▼ Microsoft Graph (2)			
User.Read	Delegated	Sign in and read user profile	No

To view and manage permissions and user consent, try [Enterprise applications](#).

Then, click on **Grant Admin consent** for Default Directory (if you don't have the permissions to click on **Grant admin consent**, please contact your IT admin to do it for you):

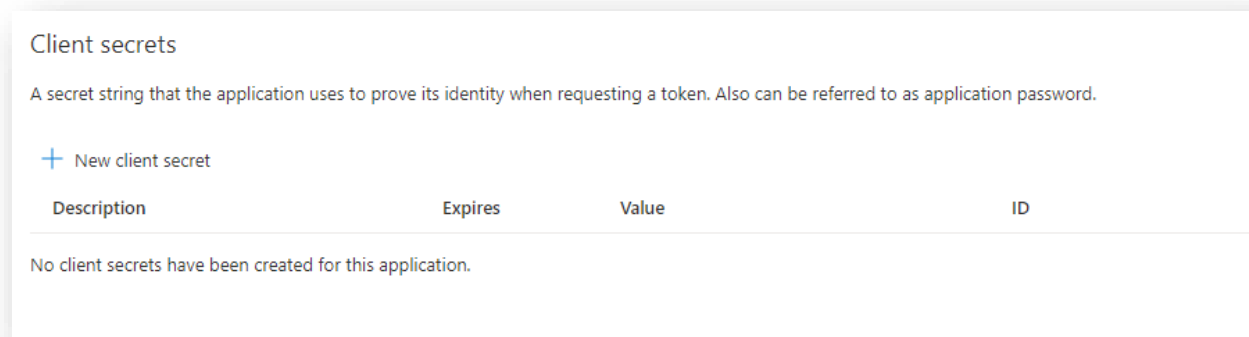
Configured permissions

Applications are authorized to call APIs when they are granted permissions by users/admins as part of the consent process. The list of configured permissions should include all the permissions the application needs. [Learn more about permissions and consent](#)

[+ Add a permission](#) [✓ Grant admin consent for Default Directory](#)

API / Permissions name	Type	Description	Admin consent req...	Status
▼ Microsoft Graph (2)				...
User.Read	Delegated	Sign in and read user profile	-	✓ Granted for Default Dire... ...

Then, take note of the client ID from the Overview tab and then go to **Certificates & Secrets** and generate a new Client Secret, take note of it.



Update the settings file from your storage account (in the clouddockitinternal folder) with the value of the previously created Azure AD Application:

```
{
  "AzureADTenant": "mytenant.onmicrosoft.com",
  "AzureADAppID": "zzzzz",
  "AzureADAppKey": "zzzzz"
}
```

Optional – Store these credentials in AWS Secrets Manager

Instead of keeping the tenant, application (client) ID and client secret in plain text in the settings file, the container can read these three values from AWS Secrets Manager. When AWS Secrets Manager is enabled, the container reads the credentials from it and only falls back to the settings file when it is not enabled or a value cannot be read. Existing setups keep working without any change.

Create one secret per credential in AWS Secrets Manager, using these names:

Secret name	Settings file field	Value
clouddockit-entra-tenant-id	AzureADTenant	Microsoft Entra tenant, e.g. mytenant.onmicrosoft.com
clouddockit-entra-client-id	AzureADAppID	Application (client) ID
clouddockit-entra-client-secret	AzureADAppKey	Application client secret

Then configure the container:

- Set the environment variable `DockerSecretsUseAwsSecretsManager` to true. If the secrets live in a specific region, also set `DockerSecretsAwsRegion` to that region.
- Grant the container's ECS task role the `secretsmanager:GetSecretValue` permission on the three secrets.

The credential values are never written to the container logs.

Step C (Optional) – Configure Clouddokit Container to support Scheduling.

Clouddokit Container supports a Scheduling Web UI that allows users to choose when they want to schedule the document generation.

To activate scheduling, you need to spin-up a new container based on the **clouddokitscheduler** image and set the appropriate settings in your settings file.

Start Clouddokit Scheduler Container

You need to follow the same procedure as you did in the previous step to spin up a new Scheduling container. You need to use the **clouddokitscheduler** image. This scheduler is basically reading the schedules files created from the UI and calling the API according to the schedule.

Here are the settings for the container:

- CPU: 1+
- RAM: 1.5GB+
- No inbound networking is required
- Outbound networking needs access to the storage account where the settings are stored and the API URL where Clouddokit is deployed.
- The following 3 environment variables are required:

Name	Value
DockerStorageCloudProvider	Specify if your Storage Account is stored in Azure, AWS or GCP. Possible values are: • Azure • GCP • AWS
DockerStorageAWSBucketName	Enter the name of the Bucket
DockerStorageAWSAccessKeyId	Enter the Access Key Id for Full control of the AWS S3 Bucket
DockerStorageAWSSecretAccessKey	Enter the Secret Access Key for Full control of the AWS S3 Bucket
DockerStorageUseAwsGov (optional)	Specify "True" if you are using Aws Gov Cloud or "False" otherwise. If the variable is not present, it's "False" by default

Set Settings in the settings file

To activate the scheduling, you need to update the settings file from your storage account (in the *clouddokitinternal* folder) to specify the URL of your Clouddokit Container:

```
{  
  "DockerUrlForSchedulingStarts": "https://myclouddokitcontainer"  
}
```

This information will be used by Clouddokit Scheduling feature to specify which Web API to call.

Step D (Optional) – Configure Clouddokit Container to support the creation of Compliance Rules, Tailored Diagrams and Settings

Clouddokit Container supports the creation of new Compliance Rules, new Tailored Diagram and new Settings.

This feature requires that you deploy an Azure Cosmos DB to save the Compliance Rules and Tailored Diagrams.


There are two steps required:

- Create (or re-use) an Azure Cosmos DB
- Add environment variables to the Clouddokit Container to specify which Azure Cosmos Database to use

Create (or re-use) an Azure Cosmos DB

From the Azure Portal, create a new Cosmos DB: (you can skip those steps if you already have a Cosmos DB that you want to reuse)

- Create a Cosmos DB



- Choose Azure Cosmos DB for NoSQL for the type

Create an Azure Cosmos DB account

Which API best suits your workload?

Azure Cosmos DB is a fully managed NoSQL and relational database service for building scalable, high performance applications. [Learn more](#)

To start, select the API to create a new account. The API selection cannot be changed after account creation.

Azure Cosmos DB for NoSQL

Azure Cosmos DB's core, or native API for working with documents. Supports fast, flexible development with familiar SQL query language and client libraries for .NET, JavaScript, Python, and Java.

[Create](#) [Learn more](#)

Azure Cosmos DB for PostgreSQL

Fully-managed relational database service for PostgreSQL with distributed query execution, powered by the Citus open source extension. Build new apps on single or multi-node clusters—with support for JSONB, geospatial, rich indexing, and high-performance scale-out.

[Create](#) [Learn more](#)

Azure Cosmos DB for Apache Cassandra

Fully managed Cassandra database service for apps written for Apache Cassandra. Recommended if you have existing Cassandra workloads that you plan to migrate to Azure Cosmos DB.

Azure Cosmos DB for Table

Fully managed database service for apps written for Azure Table storage. Recommended if you have existing Azure Table storage workloads that you plan to migrate to Azure Cosmos DB.

Once the Cosmos DB is created, you need to create a new Database named **cloudockit**:

The screenshot shows the Azure Cosmos DB Data Explorer interface. On the left is a sidebar with navigation options like Overview, Activity log, Access control (IAM), Tags, Diagnose and solve problems, Quick start, Notifications, Data Explorer, and Settings. The main area displays 'Welcome to Cosmos DB' with options to 'Start with Sample' or 'New Container'. A 'New Database' dialog box is open on the right, with the 'Database id' field containing the text 'cloudockit'. The dialog has an 'OK' button at the bottom.

Configure Clouddockit Container to use the Azure Cosmos DB

To ensure that the container can connect to the Database, you need to start the container and specify the following 2 required environment variables:


Name	Value
CosmosDb__DatabaseName	Enter the name of the Database that you have created in the previous step (clouddockit in the example)
ConnectionStrings__CosmosDb	Azure CosmosDB Connection string

Step E – Understand Clouddokit API Container

Once you have installed the Clouddokit Container, you can navigate to the Container Home Page and you will see the following screen.

It gives you the option to test the different endpoints offered by Clouddokit API.

Please note that you can do everything from command lines/scripts and not use the interface if you prefer.



For simplicity of usage, all the endpoint are POST endpoints. Not all settings are mandatory for each endpoint, and you can refer to that section to see which endpoints require which parameters.

Step F – Test your license

Activate and setup components for your license


Once you get the API Key from Cludockit team and you have the appropriate credentials for the license validation, you can check that your API Key is working by using the **/CheckLicenseStatus** endpoint.

First, navigate to the home page of the container and click on **CheckLicenseStatus** and Try it now. Then, replace the following values in the JSON that you are sending to Cludockit API:

```
{  
  "ApiKey": "API Key provided by Cludockit Team"  
}
```

Click on Execute.

You should receive the following response body:



The screenshot shows a code editor with a dark blue header labeled "Code". Below the header, the status code "200" is displayed. To the right of the status code, a JSON response body is shown in a dark gray box with white text. The JSON is as follows:

```
{  
  "data": {  
    "licenseIsValid": true  
  }  
}
```

Step G – Validate that you can authenticate to the environment that you want to scan

Once the license validation is successful, you need to test that the authentication to the environment you want to scan is working.

To do that, you need to use the `/TestAuthentication` endpoint.

First, you need to ensure that you specify the values from the above Step 2 for license validation.

Then, you need to specify the following additional values:

Name		Value
ADKCloudType		Azure/AWS/GCP depending on the platform that you want to scan.
SubscriptionID		Id/Alias of the subscription (Azure) or account (AWS) or project (GCP) that you want to scan.
(for AWS)	AWSAccessKeyId	AWS Access Key
	AWSSecretAccessKey	AWS Secret Access Key
	AWSScanRoleName	(Container only) Short role name (no ARN) assumed in each spoke account during scanning. Clouddokit builds <code>arn:aws:iam::<account>:role/<name></code> for each spoke account. Leave empty to scan using the credentials produced at login.
(for Azure)	TenantID	Tenant name of the Azure Subscription to scan
	AppClientIdForAutomation	AAD App ID for the scan
	AppClientKeyForAutomation	AAD App Key for the scan
(for GCP)	GCPServiceAccountCredentials	Content of the JSON Service Credential file
AzureStorageNameForDropOff		Do not change the name of the parameter for AWS, this is also called <code>AzureStorageNameForDropOff</code> You should specify <u>one</u> of these values: - the Azure Storage Account Name (it can be the unique Storage Account name that is in the same tenant as the subscription that you scan <i>or</i> the complete Azure Storage Account Connection String) - AWS S3 bucket - GCP Bucket where Clouddokit should store the documents generated.

Example of Payload for an AWS environment scan:

```
{
  "ApiKey": "xxxx",
  "AWSAccessKeyId": "XXXX",
  "AWSSecretAccessKey": "8PoBo+4XXXX+/k/MzQ",
  "SubscriptionID": "34XXXX2",
  "AzureStorageNameForDropOff": "XXXdockit",
  "ADKCloudType": "AWS"
}
```

Example of Payload for an Azure environment scan:

```
{
  "ApiKey": "xxxx",
  "TenantID": "X2.onmicrosoft.com",
  "AppClientIdForAutomation": "XXXXX",
  "AppClientKeyForAutomation": "mln/XXXXX=",
  "SubscriptionID": "XXX",
  "AzureStorageNameForDropOff": "XXX",
  "ADKCloudType": "Azure"
}
```

Example of Payload for an GCP environment scan:

```
{
  "ApiKey": "xxxx",
  "GCPServiceAccountCredentials": {"type":
"service_account","project_id": "cdkXXXX","private_key_id":
"XXXXX","private_key": "-----BEGIN PRIVATE KEY-----
nMIEvQIXXXXXZGy5PArVQS"n2buDji0URXCkoeWnukG9Cl0fHlP8rFK6+XXXXXX+kJm0Y
xuFOWxdbgpSln38mQyez7EK"nObnp9wP05ynOxKXJqJx0rlk="n-----END PRIVATE
KEY-----"n","client_email":
"XXXXX@cdkproject1.iam.gserviceaccount.com","client_id":
"XXXXX","auth_uri":
"https://accounts.google.com/o/oauth2/auth","token_uri":
"https://oauth2.googleapis.com/token","auth_provider_x509_cert_url"
:
"https://www.googleapis.com/oauth2/v1/certs","client_x509_cert_url":
"https://www.googleapis.com/robot/v1/metadata/x509/test-
XXXX.iam.gserviceaccount.com"},    "SubscriptionID": "XXXX",
  "AzureStorageNameForDropOff": "XXXX",
  "ADKCloudType": "GCP"
}
```

Step H – Test the document generation


Once all the tests above have been done, you can start the document generation.

To do that, you need to use the `/StartDocumentGeneration` endpoint.

First, you need to ensure that you specify the same values as the above steps for `CheckLicenseStatus` and `TestAuthentication` endpoints.

Then, you need to specify additional values based on the type of document you want to generate and which option you would like to use.

You get a list of all options from the properties list at the bottom of the screen:




The screenshot shows a web interface titled "Properties - Documentation". It features a search bar and a table with columns: Category, Title, Internal Name to use, Description, Type, and Value must be one of the following. The table lists various authentication and billing options. At the bottom, it shows "Showing 1 to 10 of 296 entries" and a pagination control with buttons for "Previous", "1", "2", "3", "4", "5", "...", "30", and "Next".

Category	Title	Internal Name to use	Description	Type	Value must be one of the following
Authentication	GCP Service Account JSON Credentials	GCPServiceAccountJSONCredentials	Specify the Service Account JSON credentials to use. This is mandatory when using the API for GCP	String	
Authentication	Tenant ID	TenantID	Specify your Azure Active Directory Tenant ID	String	
Authentication	Azure AD Application Client ID	AppClientIDForAutomation	Specify the AAD App Client ID to use for the authentication. This is mandatory when using the API for Azure	String	
Authentication	Azure AD Application Secret Key	AppClientKeyForAutomation	Specify the AAD App Secret Key to use for the authentication. This is mandatory when using the API for Azure	String	
Authentication	AWS Access Key ID	AWSAccessKeyId	Specify the AWS Access Key ID to use. This is mandatory when using the API for AWS	String	
Authentication	AWS Secret Access Key	AWSSecretAccessKey	Specify the AWS Secret Access Key to use. This is mandatory when using the API for AWS	String	
Authentication	License Code	LicenseCode	Specify your license code	String	
Billing	Dataset that contains the billing data	GCPBigQueryDataSet	Specify the name of the BigQuery Dataset that contains billing data	String	
Billing	Table that contains the billing data	GCPBigQueryTable	Specify the name of the BigQuery Table that contains the billing data.	String	
Billing	Billing Type	BillingOfferID	Specify the type of billing to use (Standard, EA or CSP)	String	

As there are many options that you can provide, we strongly advise that you use Cloudockit Website to generate the JSON file with the options.

One of the options that is particularly useful in this scenario are the `CallbackURL` and `CallbackUrlRequired` parameters that gives you the ability to be notified once document generation have been done.

When you hit Execute, you get the state URL of the current document generation:



The screenshot shows a "Server Response" with a status code of 202 and a "Response body" containing a JSON object. The JSON object includes a "stateUrl" and a "message" indicating that documentation generation was successfully started.

Code	Details
202	<pre>{ "data": { "stateUrl": "https://amazondockit.s3.us-west-2.amazonaws.com/3408512af03d8fcfe32-state.json?X-Amz-Expires=172800&X-Amz-Algorithm=AWS4-HMAC-SHA256&X-Amz-Credential=arn%3Aaws%3Aiam%3Auser%3Aamazondockit%3A%2F%2Fs3%2Faws4_request%2FX-Amz-Date=20201102T192525Z%26X-Amz-SignedHeaders=host%26X-Amz-Signature=5a37e76cf072a0266fa28ff423207f01c0fb494b7b37e802694fd5b035ba2fb4", "processId": 8420 }, "message": "Documentation generation was successfully started" }</pre>

For Payload example, you can simply re-use the previous ones.

Step I – Manage your document generation

The Clouddokit API offers two endpoints to facilitate the management of document generation.

Please note that for these endpoints, you need to specify an Admin API Key for the ApiKey value.

[/ListDocumentGeneration](#)

This will allow you to see which scans are running. It gives you the list of running processes with their Process ID and State:

Server Response	
Code	Details
202	Response body <pre>{ "data": { "processes": [{ "stateURL": "https://amazondockit.s3.us-west-2.amazonaws.com/s3/aws4_request&X-Amz-Date=20201102T192525Z&X-Amz-SignedHeader= "processID": 8420 }] } }</pre>

[/StopDocumentGeneration](#)

This endpoint is used to kill a running document generation.


Name	Value
DockerProcessToKill	Value of the process ID to kill

It will reply with a confirmation message that the process has been killed.

Server Response	
Code	Details
202	Response body <pre>{ "data": { "processKilled": true }, "message": "Process was killed" }</pre>

Annex – Deploy multiple instances of Clouddokit Container

Clouddokit can be deployed in multiple instances in scenarios like this one:




If you plan to use Clouddokit Container in a multi-pods environment, you need to configure some extra components. If you plan to use Clouddokit Containers in multiple instances with sticky session (for example App Services with a Traffic Manager), you do not need those extra components.

Here are the components that you need to configure.


Step 1 – Create / Configure Azure Key Vault

To encrypt the anti-forgery keys used by ASPNETCore, an Azure Key Vault is required. You can create a new Azure Key vault or reuse an existing one.

Once you have the Azure Key Vault, you need to create a Key named **dataprotection**




Please ensure that the Key have the following Permitted Operations (by default permissions)



Once you have done that, you need to create an Azure App Registration that will have access to this key. (you can also reuse the Azure AD App that you have created in the steps to configure Clouddockit Web UI if you prefer)

To do that, create a new App Registration (leave default settings) and note the Client ID and Client Secret as you will need that in the next steps.

Go back to the Azure Key Vault and give the Permissions to Unwrap Key / Wrap Key to the App that you just created



Step 2 – Configure Azure Redis Cache

As sessions can sprawl to multi pods, Azure Redis Cache is required to have consistent cache across all nodes.

Create a new Azure Cache for Redis (you can also reuse an existing one if you prefer) and select the Basic CO (250MB Cache) as only small elements will be cached. Ensure that you select a region that is close to the one where Cloudockit will run for performance optimization.

Once created, take note of the Redis Connection String.

Step 3 – Define the Environment Variables required to run the Cloudockit Container

In addition to the environment variables defined in the step above, you now need to add the following environment variables.

Name	Description	Example
DataProtection__EncryptionKeyUrl	URL of the key vault Key that you have created. You need to specify the Full Path to the key , not only the key vault.	https://cdkcontain erkeyvault.vault.az ure.net/keys/data protection

DataProtection__VaultClientId	Id of the Azure AD App that has privileges to Wrap / Unwrap key	760fb963-57a4-2303-1450-1b2dab513854
DataProtection__VaultSecret	Secret of the Azure AD App	SF7Q~NvuAYKF6.IB Fjdewdewd
CacheSettings__UseRedis	Set to true to use redis instead of memory cache	true
ConnectionStrings__Redis	Connection String to the Redis	cdkmultipods.redis.cache.windows.net:6380,password=xx=,ssl=True,abortConnect=False
ExternalBaseUrl	Full public URL of the container when running behind a reverse proxy or load balancer. Required for OAuth redirect URIs to resolve correctly behind a proxy.	https://clouddockit.company.com
ExternalBaseUrlProxyHeader	HTTP header whose presence signals a proxied request. When set, the ExternalBaseUrl override only applies to requests that carry this header.	X-Forwarded-Proto
TrustAllProxies	Set to True to trust all X-Forwarded-* headers. Use only when the container is behind a trusted proxy and is not directly internet-exposed. Prefer ExternalBaseUrl.	true
ForwardProxyHopLimit	Number of proxy hops to trust for X-Forwarded-For. Increase when multiple load balancers sit in front of the container.	1

For reference, here is a sample yaml file to deploy that configuration

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: clouddockit
spec:
  replicas: 4
  selector:
    matchLabels:
      app: clouddockit
  template:
    metadata:
      labels:
        app: clouddockit
    spec:
      containers:
```

```

- name: clouddockit
  image: clouddockitcontainerregistry.azurecr.io/cdk-web-linux:latest
  ports:
    - containerPort: 8080
  env:
    - name: DockerStorageCloudProvider
      value: "Azure"
    - name: "DockerStorageAzureCnxString"
      value:
"DefaultEndpointsProtocol=https;AccountName=clouddockitcontainerdebug;AccountKey=x
xxx==;EndpointSuffix=core.windows.net"
    - name: AzureADTenant
      value: "xxx.onmicrosoft.com"
    - name: AzureADAppID
      value: "9548a025-xxxx"
    - name: AzureADAppKey
      value: "W6UXfqC~xxxx"
    - name: Data__ProtectionEncryptionKeyUrl
      value:
"https://cdkcontainerkeyvault.vault.azure.net/keys/dataprotection"
    - name: DataProtection__VaultClientId
      value: "760fb9xxxx"
    - name: DataProtection__VaultSecret
      value: "VSF7xxxx"
    - name: CacheSettings__UseRedis
      value: "true"
    - name: ConnectionStrings__Redis
      value:
"cdkmultipods.redis.cache.windows.net:6380,password=xxxxk9g=,ssl=True,abortConnec
t=False"
    - name: APPINSIGHTS_CONNECTIONSTRING
      value:
"InstrumentationKey=xxx;IngestionEndpoint=xxx;LiveEndpoint=xx;ApplicationId=xxx"
    - name: TriggerDeployCount
      value: "5"
---
apiVersion: v1
kind: Service
metadata:
  name: clouddockit
spec:
  type: ClusterIP
  ports:
    - port: 80
      targetPort: 8080

```

```
selector:  
  app: clouddockit
```

Annex - AWS Container use-cases

Introduction

This annex describes the different use-cases currently available for Clouddockit Container, covering the API and the WebUI host modes, for scenarios such as scanning multiple accounts, using roles, keys, optional drop-off storage. The setups required at the Cloud platform level are listed for each use-case.

AWS ECS Container

To run Clouddockit Container in AWS, the ECS container should already be setup.

A – Scanning multiple AWS accounts using Cross-Account role

Using Clouddockit Container API, it is possible to scan multiple accounts using the Cross-Account Role functionality.

The objective is to give the possibility to a user located in the hub account to scan other AWS accounts, by assigning that user a role allowing them to access the spoke account(s) in read-only mode.

In the following example, ECS Task "EcsTask" located in the hub account with account ID "AccountA" wants to scan the resources in the spoke accounts with account IDs "AccountB", "AccountC", etc.

a. Setup in AWS Console for the role to scan multiple accounts

1. Create a role to attach to the ECS Task Definition running Clouddockit

In the hub account where the ECS is running Clouddockit, go to the IAM service, create a new role named (for example): *EcsTaskRoleForCdk*

- Permissions: Attach the following permission policies
 - *EcsCanAssumeAnyRole* (Customer Managed)
Create a new permission policy named *EcsCanAssumeAnyRole* with the following JSON:

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Effect": "Allow",
    "Action": "sts:AssumeRole",
    "Resource": "*"
  }
}
```
- Trusted entities: Attach the following Trust Relationship JSON in the role's trusted entities

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Principal": {
      "Service": "ecs-tasks.amazonaws.com"
    },
    "Action": "sts:AssumeRole"
  }]
}
```

- In the “EcsTask” Task Definition creation process, attach the created role to the ECS Task Definition (select it in the ECS Task Definition, in dropdown “Task Role”).

2. Create the Multi-Account scan role

In this step we need to define the role which is going to be used in the API to scan all the accounts.

Make sure you have the following setup in every account to scan (hub and spoke accounts).

- Create a new role named (for example): *MultiAccountScanRole*
- Permissions: Attach the following permission policies to *MultiAccountScanRole*
 - *ReadOnlyAccess* (AWS managed - job function)
 - *CdkS3BucketAccessPolicy* (Customer Managed)

Create a new permission policy named *CdkS3BucketAccessPolicy* with the following JSON:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "VisualEditor0",
    "Effect": "Allow",
    "Action": [
      "s3:GetBucketLocation",
      "s3:GetObject",
      "s3:ListBucket",
      "s3:PutBucketCORS",
      "s3:PutObject",
      "s3:DeleteObject"
    ],
    "Resource": [
      "arn:aws:s3:::*",
      "arn:aws:s3:::*/*"
    ]
  }]
}
```

- Trusted entities:
MultiAccountScanRole needs to know the calling entity, which includes the hub account ID (AccountA) of the calling role, and the role name, which is *EcsTaskRoleForCdk*.
To do that, attach the following Trust Relationship JSON in the role's Trusted entities.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Principal": {
      "AWS": [
        "arn:aws:iam::AccountA:role/EcsTaskRoleForCdk"
      ]
    },
    "Action": "sts:AssumeRole",
    "Condition": {}
  }]
}
```

b. Setup for the API Settings to use

In the Cloudockit API, in the "StartDocumentGeneration" tab, make sure to set the following settings:

- Settings defining the accounts to scan:


```
"SubscriptionId": "AccountA"
"AWSRoleToAssume": "MultiAccountScanRole"
"SelectedSubscriptionsIds": [
  "AccountA",
  "AccountB",
  "AccountC"
]
```
- To drop-off the generated documents in the S3 bucket located in "AccountA":


```
"AzureStorageNameForDropOff": "NameOfYourS3BucketInAccountA"
```
- To drop-off the generated documents in the S3 bucket located in another account (in "AccountB" for example):


```
"AzureStorageNameForDropOff":
"arn:aws:s3::AccountB:NameOfYourS3BucketInAccountB"
```
- To have the information of all scanned accounts displayed in the same document:


```
"GenerateSingleDocForAllSubscriptions": true
```

Otherwise, by default, the documents will be dropped in the storage account you specify, with a different folder for each account.

c. Setup for Web UI cross-account scanning

As of version 25, the Clouddockit Container Web UI offers two dedicated login flows for hub/spoke architecture scanning.

Option 1 — IAM user + Cross-Account Role (Access Keys)

Use this when you have a hub IAM user with static keys and a listing role in the hub account.

Field	Description
Access Key ID (required)	Access key of the hub IAM user
Secret Access Key (required)	Secret key of the hub IAM user
Cross-Account Role ARN (required)	Full ARN of the listing role in the hub account, e.g. <code>arn:aws:iam::123456789012:role/ClouddockitHubRole</code> .
Scan role name	Short name (no ARN) of the per-spoke role, e.g. <code>ClouddockitScanRole</code> . Required when the listing role does not have direct read access in the spoke accounts. Leave empty if the listing role already has cross-account read access in all spoke accounts.

Option 2 — ECS Task Role + Cross-Account Role (no keys) — Recommended

Use this when the container runs with an ECS task role and you do not want to manage static credentials.

Field	Description
Cross-Account Role ARN (required)	Full ARN of the listing role, e.g. <code>arn:aws:iam::123456789012:role/ClouddockitHubRole</code> . The ECS task role assumes this role to enumerate spoke accounts via <code>organizations:ListAccounts</code> .
Scan role name	Short name (no ARN) of the per-spoke role, e.g. <code>ClouddockitScanRole</code> . Required when the listing role does not have direct read access in the spoke accounts. Leave empty if the listing role already has cross-account read access in all spoke accounts.

IAM setup

Listing role (the Cross-Account Role ARN)

Create a policy named *ClouddockitListAccountsPolicy*

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": ["organizations:ListAccounts"],
    "Resource": "*"
  }]
}
```

Trust relationship — Option 1 (hub IAM user assumes the listing role):

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Principal": {
      "AWS": "arn:aws:iam::<HubAccountId>:user/<HubIAMUserName>"
    },
    "Action": "sts:AssumeRole"
  }]
}
```

Trust relationship — Option 2 (ECS task role assumes the listing role):

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Principal": {
      "AWS": "arn:aws:iam::<HubAccountId>:role/<EcsTaskRoleName>"
    },
    "Action": "sts:AssumeRole"
  }]
}
```

The ECS task role assumes the listing role to enumerate spoke accounts via `organizations:ListAccounts`. The listing role's trust policy must allow `sts:AssumeRole` from the task role. The scan role must exist in every spoke account under the same name and trust the listing role for `sts:AssumeRole`.

Note: The Cross-Account Role ARN is required for multi-account scanning in both options. When provided, Clouddockit uses it to enumerate spoke accounts via `organizations:ListAccounts`.

Scan role (AWSScanRoleName) — must exist in every spoke account under the same name

Permissions: attach ReadOnlyAccess (AWS managed).

Note: CdkS3BucketAccessPolicy is NOT required on the scan role. Document drop-off is handled by the hub IAM user (Option 1) or the ECS task role (Option 2), both of which already have S3 access. Only add CdkS3BucketAccessPolicy to a scan role if the drop-off bucket resides in that specific spoke account.

Trust — same for both options (listing role assumes the scan role):

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Principal": {
      "AWS": "arn:aws:iam::<HubAccountId>:role/<ListingRoleName>"
    },
    "Action": "sts:AssumeRole"
  }]
}
```

B – Using the container with the ECS Task Role instead of IAM user keys

If you don't want to use AccessKeyID/SecretAccessKey when setting up your container, you can use the ECS Task Role assigned to your container.

When creating the Task Definition, instead of setting the Environment Variables

DockerStorageAWSAccessKeyId and DockerStorageAWSSecretAccessKey

You can set the Environment Variable:

```
DockerStorageUseAwsRole = True
```

Using this option has the following impact:

- The container will access the default container storage (the storage containing the container license file) using the ECS Task role assigned to the container. This will be mainly used for validating the license, reading the user settings file, and getting the saved schedules.
- The scan will still be performed using one of the following authentication options:
 - API Settings: AccessKeyID/SecretAccessKey — scan using the hub IAM user's static credentials
 - API Settings: AWSRoleToAssume (role name, not ARN) — in Docker/ECS mode, the container's ECS task role is used as base credentials to assume the specified role. No access keys required. The role must exist in the target account under the same name with ReadOnlyAccess.
 - API Settings: AWSRoleToAssume + AccessKeyID/SecretAccessKey — the hub IAM user's keys assume the specified role.

Example payload — ECS task role with role assumption (no access keys):

```
{
  "ApiKey": "xxxx",
  "AWSRoleToAssume": "MultiAccountScanRole",
  "SubscriptionID": "123456789012",
  "AzureStorageNameForDropOff": "your-s3-bucket",
  "ADKCloudType": "AWS"
}
```

a. Setup in AWS Console for the ECS Task Role

The ECS Task Role needs to be created with the following permissions:

- Create a new role named (for example): *EcsTaskRoleContainerIdentity*
- Permissions: Attach the following permission policies to *EcsTaskRoleContainerIdentity*
 - *ReadOnlyAccess* (AWS managed - job function) (necessary if you want to scan your environment using this role)
 - *CdkS3BucketAccessPolicy* (Customer Managed)
Create a new permission policy named *CdkS3BucketAccessPolicy* with the following JSON:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "VisualEditor0",
    "Effect": "Allow",
    "Action": [
      "s3:GetBucketLocation",
      "s3:GetObject",
      "s3:ListBucket",
      "s3:PutBucketCORS",
      "s3:PutObject",
      "s3:DeleteObject"
    ],
    "Resource": [
      "arn:aws:s3:::*",
      "arn:aws:s3:::*/*"
    ]
  }]
}
```

- *EcsCanAssumeAnyRole* (Customer Managed)
Create a new permission policy named *EcsCanAssumeAnyRole* with the following JSON:

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Effect": "Allow",
    "Action": "sts:AssumeRole",
    "Resource": "*"
  }
}
```

- **Trusted entities:** attach the following Trust Relationship JSON in the role's trusted entities

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Principal": {
      "Service": "ecs-tasks.amazonaws.com"
    },
    "Action": "sts:AssumeRole"
  }]
}
```

Note: As of version 25, you can sign in using the ECS task role directly from the Cloudockit Container Web UI.

Note: The task role no longer requires the `ecs:DescribeTasks` or `ecs:DescribeTaskDefinition` permissions. Remove `AllowEcsDescribeTask` from the task role policy if it was previously configured, it is no longer needed. Cloudockit now uses `sts:GetCallerIdentity` to identify the task role, which requires no IAM grant.

C – Running a scheduled scan in AWS using ECS container

In this example, we want to run a scheduled scan from within our AWS environment. The container scheduler (the small container triggering the scheduled scans) and the container running the scan are both hosted in AWS.

a. Saving the schedule

The schedule is saved using the Web Interface (Container WebUI), we need to add the Environment Variable referring to the AWS container default S3 bucket which will now contain the saved schedules.

The 2 following steps are necessary if you are using the Container WebUI interface running in Azure to save the schedule.

1. Set the Environment Variable

- Go to a previously setup container in Azure (App service) or setup a new one (ref. `ContainerGuide_Azure`)
- In the corresponding App Service

- left panel → section **Settings**
- **Configuration**
- Click on + **New application setting**
 - Name: *DockerStorageNameForAWSScheduler*
 - Value: *NameOfYourS3Bucket* (write the name of the AWS container default S3 Bucket which contains the license file for the container running AWS)
- **Save**

2. Save the AWS Schedule

Start your Azure Web App, follow the login steps and save the schedule.

- Sign in with Azure AD
- Select "AWS", and enter the Access Key ID & Secret Access Key, log in.
- Once logged-in, set the settings you want for your schedule (make sure the IAM user with the logged-in AccessKeyID/SecretAccessKey have read access to the spoke account to scan)
- Drop-off tab: Your Storage: enter the name of any S3 bucket where you want your documents to be dropped-off (accessible with the logged-in credentials) and validate.
- Scheduling tab:
 - In the Schedules drop-down, select "Add new schedule..."
 - Enter the Schedule's Name & CRON (copy/paste CRON expression from <http://www.cronmaker.com>)
 - Enter the API Key used for the AWS Container (corresponding to the license found in the default S3 bucket)
 - Save Schedule

Your schedule file is now saved in your AWS Container default S3 bucket.

3. Run the schedule in AWS

In this final step, we are going to set the container scheduler in AWS (to trigger the scans) and have the scan executed with the container set up in AWS as well.

3.1. Set you container scheduler

- Create a new task definition/revision for the container scheduler.

In this example we will create a new Task Definition *EcsSchedulerUsingKeys* for the container scheduler, using AccessKeyID/SecretAccessKey environment variables (you can also set it with ECS Task Role instead of the keys, ref. Use-case B)

Launch Type: AWS Fargate | OS: Linux/X86_64

Task size: 1 vCPU & 2GB Memory

Task Roles

- For the ECS Task Role, attach a role (*EcsSchedulerRole*) with the following custom policies:
 - *EcsCanAssumeAnyRole* (Customer Managed) (ref. previous uses-cases)
 - *CdkS3ReadAccessPolicy* (Customer Managed)

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "VisualEditor0",
    "Effect": "Allow",
    "Action": [
      "s3:GetBucketLocation",
      "s3:GetObject",
      "s3:ListBucket",
      "s3:PutBucketCORS"
    ],
    "Resource": [
      "arn:aws:s3:::*",
      "arn:aws:s3:::*/*"
    ]
  }]
}
```

The following Trust relationships:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Principal": {
      "Service": "ecs-tasks.amazonaws.com"
    },
    "Action": "sts:AssumeRole"
  }]
}
```

- For the Task Execution Role, attach a role with the necessary permissions to access your secret manager to retrieve the secrets required to download the Clouddockit Container image (as described in ContainerGuide_AWS) and to write the container logs (in order to display any errors/warnings and to monitor the state of the schedules running).

Permissions - Policies to attach:

- *CloudWatchLogsFullAccess* (AWS Managed)
- *GetMyRegistrySecretPolicy* (Customer Managed)

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "ArnOfYourSecretFoundInSecretsManager"
    ]
  }]
}
```

The following Trust relationships:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Principal": {
      "Service": "ecs-tasks.amazonaws.com"
    },
    "Action": "sts:AssumeRole"
  }]
}
```

Container details

Set the name, and the Image URI for the scheduler container image (URI ends with *cdk-scheduler-linux:latest*), and enable Private registry authentication.

Environment Variables

Add the following Environment Variables:

- `DockerStorageCloudProvider`: AWS
- `DockerStorageAWSBucketName`: `NameOfYourContainerDefaultS3Bucket` (same as `DockerStorageNameForAWSSScheduler` set in 1. *Set the Environment Variable*)
- `DockerUrlForSchedulingStarts`: URL of the AWS container running the scan (API Host – e.g `http://57.137.30.170/`) – if this Environment Variable is not assigned in this step, it should be set in the `settings.json` file found in the S3 bucket next to the `license.json` file.

To allow the container scheduler to access the storage, you can either use IAM user access keys, **OR** the ECS Task Role to access the storage account containing the schedules:

To use the IAM user access keys:

- `DockerStorageAWSAccessKeyId` : Access Key ID of the IAM user
- `DockerStorageAWSSecretAccessKey` : Secret Access Key of the IAM user

To use the IAM Role assigned to the ECS Task (Task Role):

- `DockerStorageUseAwsRole: True`

Modify the Trust relationship to add the ECS Task role's arn in the Principal. It should look like this:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Principal": {
      "AWS": "arn:aws:iam::YourAccountId:role/EcsTaskRole",
      "Service": "ecs-tasks.amazonaws.com"
    },
    "Action": "sts:AssumeRole"
  }]
}
```

Logging

In the Logging section, Enable "Use log collection", use "Amazon CloudWatch".

The container will automatically create a log group and write the logs.

Once the task definition is created, you can deploy and run it in your cluster.

Monitor the container logs to see the status of the scan (In the running scheduler's container -> tab "Logs")

Annex – Troubleshooting

Here are resolutions to common cases and how you can help find errors in Clouddokit Container.

- If you activate Clouddokit Container Web UI and noticed that in the upper right corner you have a Welcome message without your name, please check the AAD Credentials in the settings file
- If you are using Private endpoint for your App Service and Storage, please ensure that you activate vNET integration so that the App Service can communicate with the Storage Account
- You can specify an environment variable in your container named `AppInsightsConnectionString` that contains an Azure Application Insights Connection String so that you can see the logs.
- You can use the `-logs.txt` file in the storage that you have specified to see what is happening during document generation.
- If you get an error when the document generation starts, please ensure that you have Write privileges to your storage account
- If you see the message that the document generation is starting but do not see any progress, please verify that you have a CORS rule for GET Verb and origin that is your Clouddokit container website (should be done automatically).
- If you get an exception when starting the container that says “APPCMD failed with error code 87”, check that the variables that you are providing do not contain quotes.