Implementation Guide

Serverless Image Handler

Serverless Image Handler: Implementation Guide

Table of Contents

Serverless architecture for cost-effective image processing

Publication date: *June 2017 ([last update](): June 2024)*

The Serverless Image Handler solution helps you embed images on your websites and mobile applications to drive user engagement. It uses the `sharp` Node.js library to provide high-speed image processing without sacrificing image quality. To minimize your costs of image optimization, manipulation, and processing, this solution automates version control and provides flexible storage and compute options for file reprocessing.

This solution automatically deploys and configures a serverless architecture optimized for dynamic image manipulation. Images can be rendered and returned spontaneously. For example, you can automate resizing of an image based on different screen sizes by adding code on your website that leverages this solution. This helps you adapt your website's presentation to meet your users' different modes of viewing. This solution uses [Amazon CloudFront]() for global content delivery and [Amazon Simple Storage Service]() (Amazon S3) for reliable and durable cloud storage.

This implementation guide provides an overview of the Serverless Image Handler solution, its reference architecture and components, considerations for planning the deployment, configuration steps for deploying the solution to the Amazon Web Services (AWS) Cloud.

The intended audience for implementing this solution in their environment includes solution architects, business decision makers, DevOps engineers, data scientists, and cloud professionals.

Use this navigation table to quickly find answers to these questions:

If you want to . . .	Read . . .
Know the cost for running this solution. The estimated cost for running this solution in the US East (N. Virginia) Region is **approximately USD $15.44 per month** for 100,000 new images.	[Cost]()
Understand the security considerations for this solution.	[Security]()

If you want to . . .	Read . . .
Know how to plan for quotas for this solution.	Quotas
Know which AWS Regions support this solution.	Supported AWS Regions
View or download the AWS CloudFormation template included in this solution to automatically deploy the infrastructure resources (the "stack") for this solution.	AWS CloudFormation template
Access the source code and optionally use the AWS Cloud Development Kit (AWS CDK) to deploy the solution.	GitHub repository

Features and benefits

This solution provides the following features:

Dynamic content delivery

Automatically modify images based on users' devices and screen sizes.

Content moderation

Use Amazon Rekognition to automatically detect and blur inappropriate user-uploaded images.

Smart cropping

Use Amazon Rekognition to crop images using facial recognition.

Low-cost image storage

Save on image storage costs by generating modified images at runtime, and caching generated images in CloudFront.

Integration with Service Catalog AppRegistry and Application Manager, a capability of AWS Systems Manager

This solution includes a Service Catalog AppRegistry resource to register the solution's AWS CloudFormation template and its underlying resources as an application in both AppRegistry and Application Manager. With this integration, you can centrally manage the solution's resources.

Use cases

Drive user engagement

Improve engagement with your mobile application or website by maintaining high-quality images that adjust for device screen size.

Improve user and brand safety

Automatically detect and blur inappropriate user-uploaded images with machine learning trained to recognize pre-defined and user-defined categories .

Concepts and definitions

This section describes key concepts and defines terminology specific to this solution:

cross-origin resource sharing (CORS)

Defines a way for client web applications that are loaded in one domain to interact with resources in a different domain.

fallback image

Image that you set to show when the intended image doesn't load.

> **ⓘ Note**
>
> For a general reference of AWS terms, see the AWS Glossary.

Architecture overview

This section provides a reference implementation architecture diagram for the components deployed with this solution.

Architecture diagram

Deploying this solution with the default parameters deploys the following components in your AWS account.

> ⚠️ **Important**
>
> This solution is intended for customers with public applications who want to provide an option to dynamically change or manipulate their public images. Because of these public requirements, this template creates a publicly accessible, unauthenticated CloudFront distribution and Amazon API Gateway endpoint in your account, allowing anyone to access it. For more information on API Gateway authorization, refer to the Security section.
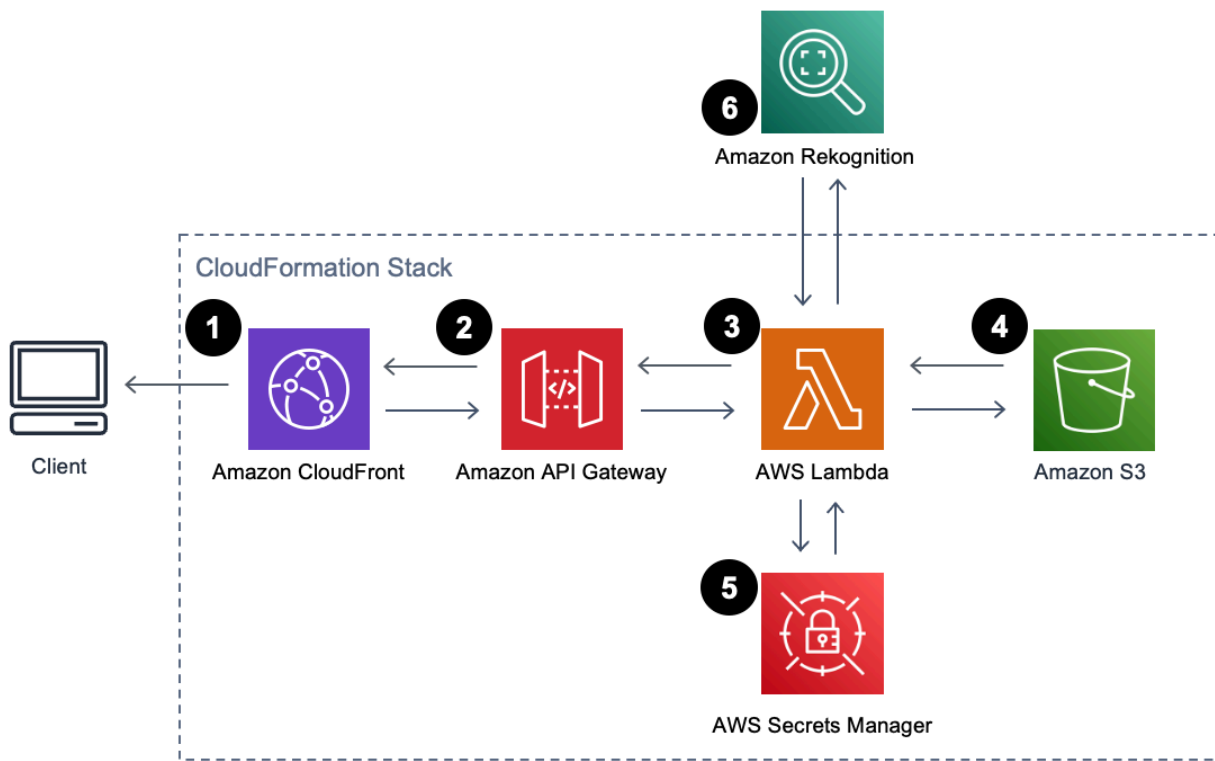
Serverless Image Handler architecture on AWS

> ⓘ **Note**
>
> AWS CloudFormation resources are created from [AWS Cloud Development Kit](link) (AWS CDK) constructs.

The high-level process flow for the solution components deployed with the AWS CloudFormation template is as follows:

1. An [Amazon CloudFront](link) distribution provides a caching layer to reduce the cost of image processing and the latency of subsequent image delivery. The CloudFront domain name provides cached access to the image handler application programming interface (API).

2. [Amazon API Gateway](link) provides endpoint resources and initiate the [AWS Lambda](link) function.

3. A Lambda function retrieves the image from a customer's existing [Amazon S3](link) bucket and uses `sharp` to return a modified version of the image to the API Gateway.

4. A solution-created S3 bucket provides log storage, separate from your customer-created S3 bucket for storing images. If you enter `Yes` (default entry) for **the Deploy Demo UI** template parameter, the solution deploys another S3 bucket for storing the optional [demo user interface](link) (UI).

5. (Optional) If you enter `Yes` for the **Enable Signature** template parameter, the Lambda function retrieves the secret value from your existing [AWS Secrets Manager](link) secret to validate the signature. For more information, see [Launch the stack](link).

6. (Optional) If you use the [smart crop](link) or [content moderation](link) features, the Lambda function calls [Amazon Rekognition](link) to analyze your image and returns the results.

Architecture details

This section describes the components and AWS services that make up this solution and the architecture details on how these components work together.

Demo UI

This solution optionally deploys a demo UI into your account to demonstrate the basic features of the solution. You can use the UI to interact directly with your new image handler API endpoint, using image files that already exist in your account.

This solution's template contains a **Deploy Demo UI** parameter that's activated (set to Yes) by default. If activated, this option deploys an additional Amazon S3 bucket and associated CloudFront distribution into your account.

Smart cropping

You can use this image request option to crop images using the facial recognition capabilities of Amazon Rekognition. To generate a cropped image, a Lambda function sends requests to Amazon Rekognition to identify faces in images and calculate crop areas.

> ⓘ **Note**
>
> Amazon Rekognition supports only JPEG and PNG file formats for smart cropping. When using the Amazon Rekognition features with an image that isn't JPEG or PNG, the solution automatically converts the image to PNG for use with Amazon Rekognition, then converts it back to the original format.

Content moderation

You can use this image request option to detect and blur inappropriate images. To detect an inappropriate image, a Lambda function sends requests to Amazon Rekognition to identify inappropriate content.

> ⓘ **Note**
>
> Amazon Rekognition supports only JPEG and PNG file formats for content moderation. When using the Amazon Rekognition features with an image that isn't JPEG or PNG, the solution automatically converts the image to PNG for use with Amazon Rekognition, then converts it back to the original format.

Cross-origin resource sharing

This solution's template contains two parameters that activate Cross-origin resource sharing (CORS) for your image handler API: **CorsEnabledParameter** and **CorsOriginParameter**. CORS defines how client web applications loaded in one domain can interact with resources in a different domain. You can activate CORS for your image handler API to make requests to your image handler API from outside the domain space of the API.

For example, if you have a public web application hosted on either a custom domain or a cloud domain outside of AWS, you can activate CORS to fetch original or modified images from the image handler API.

> ⓘ **Note**
>
> If you want to change your CORS configuration after deployment, you can activate or deactivate CORS by editing the `CORS_ENABLED (Yes/No)` and `CORS_ORIGIN` environment variables of the Lambda image handler function. See Using AWS Lambda environment variables in the *AWS Lambda Developer Guide* for more information.

Image URL signature

This solution's template contains three parameters that are required for the image URL signature functionality: **EnableSignatureParameter**, **SecretsManagerSecretParameter**, and **SecretsManagerKeyParameter**. To activate this feature:

- Set the **EnableSignatureParameter** parameter to Yes
- Set the **SecretsManagerSecretParameter** and **SecretsManagerKeyParameter** parameters to a valid secret and key that you originally created in Secrets Manager

> ⚠ **Important**
>
> You are responsible for creating the Secrets Manager secret and key. For more information about Secrets Manager secret creation, refer to [Create and manage secrets with AWS Secrets Manager](#) in the *AWS Secrets Manager User Guide*.

When you activate this feature, the image handler AWS Lambda function checks for a valid signature in the image request. If the signature doesn't match, the solution returns an error message. When activating the image URL signature, you must provide the `signature` query string to your URL. For example, you can create the signature using the following Node.js code:

```
const secret = '<YOUR_SECRET_VALUE_IN_SECRETS_MANAGER>';
const path = '/<YOUR_PATH>'; // Add the first '/' to path.
const signature = crypto.createHmac('sha256', secret).update(path).digest('hex');
```

You can request your image using the image URL signature:

```
https://<distributionName>.cloudfront.net/<YOUR_PATH>?signature=<YOUR_SIGNATURE>
```

> ⓘ **Note**
>
> If you update your existing solution deployment and activate the image URL signature, the updated stack will no longer be compatible with the existing URLs. You must update your application to provide the correct signature query string to your URLs. To update the solution stack, refer to [Update the solution](#).

Default fallback image

This solution provides a default [fallback image](#) feature that returns the specified fallback image as a result of errors occur during processing, rather than a JSON object error message. This solution's template contains three parameters that are required for the default fallback image feature: **EnableDefaultFallbackImageParameter**, **FallbackImageS3BucketParameter**, and **FallbackImageS3KeyParameters**.

By default, this feature is deactivated. To activate this feature:

> **ⓘ Note**
>
> Before activating this feature, if you use an S3 bucket policy in the fallback image S3 bucket, you must edit the bucket policy to allow the `CustomResourceFunction` and `ImageHandlerFunction` AWS Lambda functions to get the default fallback image object. For more information, see [Adding a bucket policy by using the Amazon S3 console](#).

- Set the **EnableDefaultFallbackImageParameter** parameter to Yes
- Set the **FallbackImageS3BucketParameter** and **FallbackImageS3KeyParameter** parameters to a valid S3 bucket and object key

AWS services in this solution

AWS service	Description
[Amazon CloudFront](#)	**Core.** Provides a caching layer to reduce latency and the cost of image processing for subsequent identical requests.
[AWS Lambda](#)	**Core.** Runs functions to retrieve, modify, and invoke other services to analyze images. Also runs a function to support URL signature validation.
[Amazon S3](#)	**Core.** Stores images, logs, and a demo UI.
[Amazon API Gateway](#)	**Supporting.** Provides API endpoints to invoke Lambda functions.
[AWS CDK](#)	**Supporting.** Provides infrastructure as code constructs to generate the solution's underlying CloudFormation templates.
[AWS CloudFormation](#)	**Supporting.** Deploys the solution's underlying AWS resources.

AWS service	Description
AWS Identity and Access Management (IAM)	**Supporting.** Allows for fine-grained access permissions.
Amazon Rekognition	**Optional.** Uses machine learning (ML) to analyze images.
AWS Secrets Manager	**Optional.** Manages secrets to support URL signatures.

Plan your deployment

This section describes the [cost](#), [security](#), [quotas](#), and other considerations before deploying the solution.

Supported AWS Regions

This solution uses AWS services that aren't available in all AWS Regions. You must launch this solution in an AWS Region where these services are available. For the current availability of AWS services by Region, see the [AWS Regional Services List](#).

This solution is available in the following AWS Regions:

Region name	
US East (Ohio)	Canada (Central)
US East (N. Virginia)	China (Beijing)
US West (Northern California)	China (Ningxia)
US West (Oregon)	Europe (Frankfurt)
Africa (Cape Town)	Europe (Ireland)
Asia Pacific (Hong Kong)	Europe (London)
Asia Pacific (Mumbai)	Europe (Milan)
Asia Pacific (Seoul)	Europe (Paris)
Asia Pacific (Singapore)	Europe (Stockholm)
Asia Pacific (Sydney)	Middle East (Bahrain)
Asia Pacific (Tokyo)	South America (São Paulo)

Opt-in Regions

An opt-in Region is an AWS Region that's deactivated by default. You can activate opt-in Regions can be activated in the AWS console. For additional information about opt-in Regions and how to activate them, refer to Managing AWS Regions in the *AWS General Reference guide*.

This solution supports four opt-in Regions:

- Asia Pacific (Hong Kong)

- Middle East (Bahrain)

- Africa (Cape Town)

- Europe (Milan)

  When launched in an opt-in Region, this solution creates an S3 logging bucket for CloudFront in the US East (N. Virginia) Region. This is because CloudFront doesn't deliver access logs to buckets in the supported opt-in Regions. For more information about S3 buckets, refer to Choosing an Amazon S3 bucket for your standard logs in the *Amazon CloudFront Developer Guide*.

To deploy in an opt-in Region, the S3 bucket that you provide for the **Source Buckets** parameter must be in the same Region where launch the CloudFormation template.

Cost

You are responsible for the cost of the AWS services used while running this solution. As of this revision, the cost for running this solution with the default settings in the US East (N. Virginia) Region is approximately **$15.44 per month** for 100,000 new images, **$150.83 per month** for 1,000,000 new images, and **$752.57 per month** for 5,000,000 new images (refer to the Sample cost table for the cost breakdown).

We recommend creating a budget through AWS Cost Explorer to help manage costs. Prices are subject to change. For full details, see the pricing webpage for each AWS service used in this solution.

Sample cost table

The following table provides a sample cost breakdown for deploying this solution with the default parameters in the US East (N. Virginia) Region for one month.

AWS service	Dimensions	Cost [USD] (100,000 new images)	Cost [USD] (1,000,000 new images)	Cost [USD] (5,000,000 new images)
Amazon API Gateway		$0.35	$3.50	$17.50
AWS Lambda	2 seconds processing time per image	$3.35	$33.53	$167.67
Amazon CloudFront	process 1 MB images per request	$8.50	$85.00	$425.00
Amazon S3	store 1 MB images	$2.34	$23.40	$117.00
AWS Secrets Manager*		$0.90	$5.40	$25.40
Total		**$15.44**	**$150.83**	**$752.57**

* The cost for AWS Secrets Manager is incurred only when the image URL signature feature is activated.

Demo UI

If you choose to deploy the demo UI, the solution automatically deploys an additional CloudFront distribution and S3 bucket for storing the static website assets in your account. You are responsible for the incurred variable charges from these services. For more information, see Amazon S3 pricing.

Image modification and analysis

This cost estimate doesn't account for Amazon S3 PUT and GET requests, which can vary because modified images are cached in CloudFront, and because certain use cases require special-use capabilities such as smart cropping and content moderation with Amazon Rekognition. Using

Amazon Rekognition features may incur additional charges. For more information, see Amazon Rekognition pricing.

There is no additional cost for using `sharp`, which is an open source library.

Security

When you build systems on AWS infrastructure, security responsibilities are shared between you and AWS. This shared responsibility model reduces your operational burden because AWS operates, manages, and controls the components including the host operating system, the virtualization layer, and the physical security of the facilities in which the services operate. For more information about AWS security, visit AWS Cloud Security.

> ⚠️ **Important**
>
> This solution creates CloudFront and API Gateway resources that are publicly accessible. Be aware that while this is likely appropriate for publicly facing websites, it might not be appropriate for all customer use cases for this solution.
> AWS offers several options for end-to-end security, such as AWS Identity and Access Management (IAM), Amazon Cognito user pools, AWS Certificate Manager, and CloudFront signed URLs. For private image handling use cases, AWS recommends using CloudFront signed URLs and implementing an API Gateway Lambda authorizer with CloudFront to secure your stack.

Demo UI

This solution optionally deploys a demo UI as a static website hosted in an S3 bucket. To help reduce latency and improve security, this solution includes a CloudFront distribution with an origin access identity, which is a CloudFront user that helps restrict access to the solution's website S3 bucket contents. For more information, refer to Restricting access to an Amazon S3 origin in the *Amazon CloudFront Developer Guide*.

IAM roles

IAM roles allow customers to assign granular access policies and permissions to services and users on the AWS Cloud. This solution creates IAM roles that grant the solution's Lambda functions access to create Regional resources.

Quotas

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. Make sure you have sufficient quota for each of the services implemented in this solution. For more information, see AWS service quotas.

Use the following links to go to the page for that service. To view the service quotas for all AWS services in the documentation without switching pages, view the information in the Service endpoints and quotas page in the PDF instead.

AWS CloudFormation quotas

Your AWS account has CloudFormation quotas that you should consider when launching the stack for this solution. By understanding these quotas, you can avoid limitation errors that can prevent you from deploying this solution successfully. For more information, see AWS CloudFormation quotas in the in the *AWS CloudFormation User Guide*.

AWS Lambda quotas

Lambda has a 6 MB invocation payload request and response limit. For information about Lambda quotas, including the amount of compute and storage resources that you can use to run and store functions, refer to Lambda quotas in the *AWS Lambda Developer Guide*.

Amazon API Gateway quotas

API Gateway sets the maximum integration timeout at 30 seconds for all integration types, including Lambda. Processing large image files can result in a timeout error due to the maximum integration timeout being exceeded. For information about API Gateway quotas, refer to Amazon API Gateway quotas and important notes in the *Amazon API Gateway Developer Guide*.

Deploy the solution

This solution uses [CloudFormation templates and stacks](#) to automate its deployment. The CloudFormation template specifies the AWS resources included in this solution and their properties. The CloudFormation stack provisions the resources that are described in the templates.

Before you launch the solution, review the [cost](#), [architecture](#), [network security](#), and other considerations discussed earlier in this guide.

> ⚠️ **Important**
>
> This solution includes an option to send anonymized operational metrics to AWS. We use this data to better understand how customers use this solution and related services and products. AWS owns the data gathered though this survey. Data collection is subject to the [AWS Privacy Notice](#).
> To opt out of this feature, download the template, modify the AWS CloudFormation mapping section, and then use the AWS CloudFormation console to upload your updated template and deploy the solution. For more information, see the [Anonymized data collection](#) section of this guide.

AWS CloudFormation template

You can download the CloudFormation template for this solution before deploying it.

[View template] **serverl**

image-handler.template – Use this template to launch the solution and all associated components. The default configuration deploys CloudFront, API Gateway, Lambda, and Amazon S3. You can customize the template to meet your specific needs.

> ⓘ **Note**
>
> CloudFormation resources are created from AWS CDK constructs.

Before you launch the solution's AWS CloudFormation template, you must specify an S3 bucket in the **Source Buckets** template parameter. Use this S3 bucket to store the images that you want

to manipulate. If you have multiple image source S3 buckets, you can specify them as comma-separated values. For lower latency, use an S3 bucket in the same AWS Region where you launch your CloudFormation template.

> ⓘ **Note**
>
> If you are launching from a supported opt-in Region, the source S3 bucket you created and provided as the **Source Buckets** template parameter must be in the same Region where you're launching the CloudFormation template.

We recommend deploying the optional demo UI when you first deploy the solution to test the solution's functionality. For more information, refer to Use the demo UI.

> ⓘ **Note**
>
> If you have previously deployed this solution, see Update the solution for update instructions.
> Serverless Image Handler version 6.0 and newer include significant changes, and you can't update the solution from versions before 6.0 to version 6.0 or later. To use version 6.0 or later, launch a new stack using version 6.x of the CloudFormation template and uninstall your previous version of this solution.

Launch the stack

Follow the step-by-step instructions in this section to configure and deploy the solution into your account.

Time to deploy: Approximately 15 minutes

1. Sign in to the AWS Management Console and select the button to launch the `serverless-image-handler` AWS CloudFormation template.

Launch solution

2. The template launches in the US East (N. Virginia) Region by default. To launch the solution in a different AWS Region, use the Region selector in the console navigation bar. For a list of which AWS Regions support this solution, see Supported AWS Regions.

3. On the **Create stack** page, verify that the correct template URL is in the **Amazon S3 URL** text box and choose **Next**.

4. On the **Specify stack details** page, assign a name to your solution stack. For information about naming character limitations, see IAM and AWS STS quotas in the *AWS Identity and Access Management User Guide*.

5. Under **Parameters**, review the parameters for this solution template and modify them as necessary. This solution uses the following default values.

Parameter	Default	Description
CORS Enabled	No	Choose whether to activate CORS. For information about this parameter, refer to Cross-origin resource sharing (CORS).
CORS Origin	*	This value is returned by the API in the **Access-Control-Allow-Origin** header. An asterisk (*) value supports any origin. We recommend specifying a specific origin (such as `http://<example>.<domain>`) to restrict cross-site access to your API. ⓘ **Note** This value is ignored if the **CORS Enabled**

Parameter	Default	Description
		parameter is set to No.
Source Buckets	*<Requires input>*	Specifies the S3 bucket (or buckets) in your account that contain(s) the images that you manipulate. To specify multiple buckets, separate them by commas.
Deploy Demo UI	Yes	The demo UI that deploys to the Demo S3 bucket. For more information refer to Use the demo UI.
Log Retention Period	1	Specifies the number of days to retain Lambda log data in CloudWatch logs.
Enable Signature	No	Choose whether to activate the image URL signature feature. For information about this feature, refer to Image URL signature.

Parameter	Default	Description
SecretsManager Secret	*<Optional input>*	Define the Secrets Manager secret name that contains the secret key for the image URL signature. ⓘ **Note** This value is ignored if the **Enable Signature** parameter is set to No.
SecretsManager Key	*<Optional input>*	Define the Secrets Manager secret key that contains the secret value to create the image URL signature. ⓘ **Note** This value is ignored if the **Enable Signature** parameter is set to No.
Enable Default Fallback Image	No	Choose whether to activate the default fallback image feature. For information about this feature, refer to [Default fallback image](#).

Parameter	Default	Description
Fallback Image S3 Bucket	*<Optional input>*	Specify the S3 bucket which contains the default fallback image. ⓘ **Note** This value is ignored if the **Enable Default Fallback Image** parameter is set to No.
Fallback Image S3 Key	*<Optional input>*	Specify the default fallback image S3 object key, including prefix. See Creating object key names for more information. ⓘ **Note** This value is ignored if the **Enable Default Fallback Image** parameter is set to No.
AutoWebP	No	Choose whether to automatically accept WebP image formats.

Parameter	Default	Description
CloudFront PriceClass	PriceClass_All	The CloudFront price class to use. For more information, refer to Choosing the price class for a CloudFront distribution in the *Amazon CloudFront Developer Guide*.

6. Choose **Next**.

7. On the **Configure stack options** page, choose **Next**.

8. On the **Review and create** page, review and confirm the settings. Select the box acknowledging that the template creates IAM resources.

9. Choose **Submit** to deploy the stack.

   You can view the status of the stack in the AWS CloudFormation console in the **Status** column. You should receive a CREATE_COMPLETE status in approximately 15 minutes.

Monitor the solution with AppRegistry

The solution includes a Service Catalog AppRegistry resource to register the CloudFormation template and underlying resources as an application in both Service Catalog AppRegistry and AWS Systems Manager Application Manager.

AWS Systems Manager Application Manager gives you an application-level view into this solution and its resources so that you can:

- Monitor its resources, costs for the deployed resources across stacks and AWS accounts, and logs associated with this solution from a central location.
- View operations data for the resources of this solution in the context of an application. For example, deployment status, CloudWatch alarms, resource configurations, and operational issues.

The following figure depicts an example of the application view for the solution stack in Application Manager.

Solution stack in Application Manager

Activate CloudWatch Application Insights

1. Sign in to the [Systems Manager console](#).

2. In the navigation pane, choose **Application Manager**.

3. In **Applications**, search for the application name for this solution and select it.

   The application name will have **App Registry** in the **Application Source** column, and will have a combination of the solution name, Region, account ID, or stack name.

4. In the **Components** tree, choose the application stack you want to activate.

5. In the **Monitoring** tab, in **Application Insights**, select **Auto-configure Application Insights**.

Monitoring for your applications is now activated and the following status box appears:

Confirm cost tags associated with the solution

After you activate cost allocation tags associated with the solution, you must confirm the cost allocation tags to see the costs for this solution. To confirm cost allocation tags:

1. Sign in to the Systems Manager console.

2. In the navigation pane, choose **Application Manager**.

3. In **Applications**, choose the application name for this solution and select it.

   The application name will have **App Registry** in the **Application Source** column, and will have a combination of the solution name, Region, account ID, or stack name.

4. In the **Overview** tab, in **Cost**, select **Add user tag**.
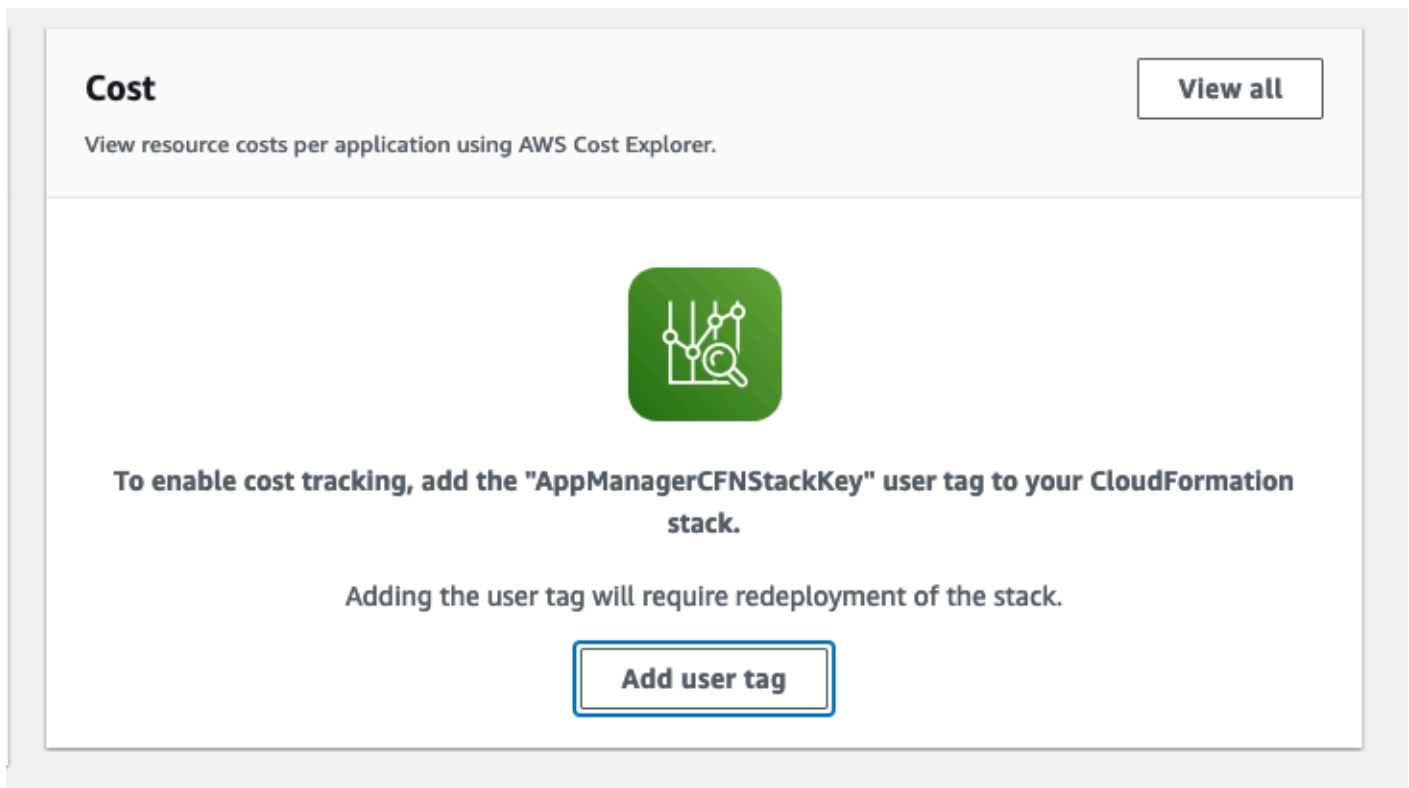
5. On the **Add user tag** page, enter `confirm`, then select **Add user tag**.

The activation process can take up to 24 hours to complete and the tag data to appear.

Activate cost allocation tags associated with the solution

After you activate Cost Explorer, you must activate the cost allocation tags associated with this solution to see the costs for this solution. The cost allocation tags can only be activated from the management account for the organization. To activate cost allocation tags:

1. Sign in to the [AWS Billing and Cost Management and Cost Management console](#).
2. In the navigation pane, select **Cost Allocation Tags**.
3. On the **Cost allocation tags** page, filter for the AppManagerCFNStackKey tag, then select the tag from the results shown.
4. Choose **Activate**.

AWS Cost Explorer

You can see the overview of the costs associated with the application and application components within the Application Manager console through integration with AWS Cost Explorer, which must be first activated. Cost Explorer helps you manage costs by providing a view of your AWS resource costs and usage over time. To activate Cost Explorer for the solution:

1. Sign in to the [AWS Cost Management console](#).
2. In the navigation pane, select **Cost Explorer** to view the solution's costs and usage over time.

Update the solution

If you have previously deployed the solution, follow this procedure to update the CloudFormation stack to get the latest version of the solution's framework.

> ⚠️ **Important**
>
> Serverless Image Handler version 6.0 and newer include significant changes, and you can't update the solution from versions prior to 6.0 to version 6.0 or later. To use version 6.0 or later, launch a new stack using version 6.x of the CloudFormation template and uninstall your previous version of this solution.

1. Sign in to the AWS CloudFormation console, select your existing Serverless Image Handler CloudFormation stack, and select **Update**.

2. Select **Replace current template**.

3. Under **Specify template**:

   a. Select **Amazon S3 URL**.

   b. Copy the link of the `serverless-image-handler.template` the section called "AWS CloudFormation template".

   c. Paste the link in the **Amazon S3 URL** box.

   d. Verify that the correct template URL shows in the **Amazon S3 URL** text box, and choose **Next**. Choose **Next** again.

4. Under **Parameters**, review the parameters for the template and modify them as necessary. For details about the parameters, see Deployment process overview.

5. Choose **Next**.

6. On the **Configure stack options** page, choose **Next**.

7. On the **Review** page, review and confirm the settings. Select the box acknowledging that the template creates IAM resources.

8. Choose **View change set** and verify the changes.

9. Choose **Update stack** to deploy the stack.

You can view the status of the stack in the AWS CloudFormation console in the **Status** column. You should receive an UPDATE_COMPLETE status in approximately 15 minutes.

Backward compatibility

This solution is compatible with legacy image request formats, including the Thumbor and Custom (with rewrite function) formats from previous versions of this solution. If you are using a previous version of this solution (version 3.x and earlier) and have image requests formatted for use with that version, review the following note to ensure minimal breaking changes or parities.

> **ⓘ Note**
>
> Legacy requests (Thumbor and custom) will source images from the first bucket in the SOURCE_BUCKETS environment variable by default. To use a different bucket, you can use the s3:BucketName tag in your request or you can adjust which bucket is first in the environment variables section of your image handler Lambda function. See Using AWS Lambda environment variables in the *AWS Lambda Developer Guide* for more information.

Thumbor compatibility

You can specify Thumbor image requests as you normally would, with filters and other relevant properties added on as suffixes to the default CloudFront **ApiEndpoint**. For more information about using Thumbor, see List of supported Thumbor filters.

> **ⓘ Note**
>
> Serverless Image Handler includes a Thumbor-style interface in the API; however, those requests are mapped to comparable Sharp library calls, and might not include all available Thumbor filters. For more information about available Thumbor-style filters, see List of supported Thumbor filters.

Custom compatibility

You can specify custom image requests that used the version 3.x and earlier solution versions' rewrite feature as you normally would. First, you must update the REWRITE_MATCH_PATTERN and REWRITE_SUBSTITUTION environment variables for your image handler function with the appropriate (JavaScript/ECMAScript-compatible) regular expressions and strings. For example:

```
https://<distName>.cloudfront.net/<customRequestHere>
```

For more information about using custom image requests, see [Custom image requests](#).

Troubleshooting

If you need help with this solution, contact AWS Support to open a support case for this solution.

Contact AWS Support

If you have [AWS Developer Support](), [AWS Business Support](), or [AWS Enterprise Support](), you can use the Support Center to get expert assistance with this solution. The following sections provide instructions.

Create case

1. Sign in to [Support Center]().
2. Choose **Create case**.

How can we help?

1. Choose **Technical**.
2. For **Service**, select **Solutions**.
3. For **Category**, select **Other Solutions**.
4. For **Severity**, select the option that best matches your use case.
5. When you enter the **Service**, **Category**, and **Severity**, the interface populates links to common troubleshooting questions. If you can't resolve your question with these links, choose **Next step: Additional information**.

Additional information

1. For **Subject**, enter text summarizing your question or issue.
2. For **Description**, describe the issue in detail.
3. Choose **Attach files**.
4. Attach the information that AWS Support needs to process the request.

Help us resolve your case faster

1. Enter the requested information.

2. Choose **Next step: Solve now or contact us**.

Solve now or contact us

1. Review the **Solve now** solutions.

2. If you can't resolve your issue with these solutions, choose **Contact us**, enter the requested information, and choose **Submit**.

Uninstall the solution

You can uninstall the solution from the [AWS Management Console](#) or by using the [AWS Command Line Interface](#) (AWS CLI). You must manually delete the S3 buckets created by this solution. AWS solutions don't automatically delete these resources in case you have stored data to retain.

Using the AWS Management Console

1. Sign in to the [AWS CloudFormation console](#).

2. On the **Stacks** page, select this solution's installation stack.

3. Choose **Delete**.

Using AWS Command Line Interface

Determine whether the AWS CLI is available in your environment. For installation instructions, see [What Is the AWS Command Line Interface?](#) in the *AWS CLI User Guide*. After confirming that the AWS CLI is available, run the following command.

```
$ aws cloudformation delete-stack --stack-name <installation-stack-name>
```

Deleting the Amazon S3 buckets

This solution is configured to retain the solution-created S3 buckets if you decide to delete the AWS CloudFormation stack to prevent accidental data loss. After uninstalling the solution, you can manually delete this S3 bucket if you don't need to retain the data. Follow these steps to delete the Amazon S3 buckets.

1. Sign in to the [Amazon S3 console](#).

2. Choose **Buckets** from the left navigation pane.

3. Locate the *<stack-name>* S3 buckets.

4. Select the S3 bucket and choose **Delete**.

To delete the S3 bucket using AWS CLI, run the following command:

Using the AWS Management Console   32

```
$ aws s3 rb s3://<bucket-name> --force
```

Alternatively, you can configure the CloudFormation template to delete the Amazon
S3 bucket automatically. Before deleting the stack, change the deletion behavior in the
CloudFormation `DeletionPolicy attribute`.

> ⓘ **Note**
>
> Neither of these methods deletes the source bucket you created and provided as a
> parameter to the CloudFormation template.

Use the solution

This section provides a user guide for utilizing the AWS solution.

Use the demo UI

The solution provides an optional demo UI that you can deploy into your AWS account to display basic capability and functionality. With this UI, you can interact directly with the new image handler using images from the specified Amazon S3 buckets in your account.

Serverless Image Handler demo UI

Follow this procedure to experiment with the supported image editing features, preview the results, and create example URLs that you can use in your applications:

1. Sign in to the [AWS CloudFormation console](#).
2. Select the solution's installation stack.
3. Choose the **Outputs** tab, and then select value for the **DemoUrl** key. The Serverless Image Handler Demo UI opens in your browser.
4. In the **Image Source** card, perform the following actions:

   a. Specify a bucket name to use for the demo. The bucket you specify must be listed in the SOURCE_BUCKETS environment variable of the AWS Lambda function.

b. Specify an image key to use for the demo. You must include the file extension in the key.

5. Select **Import**. The original image appears in the **Original Image** card.

6. In the **Editor** card, adjust the image settings, and select **Preview** to generate the modified image. You can select **Reset** to revert the settings back to their original values.

> ### ⓘ Note
>
> The Serverless Image Handler demo UI offers a limited set of image edits and doesn't include the full scope of capabilities offered by the Image Handler API and the image URL signature. We recommended using your own [frontend application](#) for image modification.

Use the solution with a frontend application

In your frontend application, you can access both the original and modified images by creating an image request object, stringifying and encoding that object, and appending it to the API call. Follow these steps to retrieve your API endpoint for the solution:

1. Sign in to the [AWS CloudFormation console](#).

2. On the **Stacks** page, select this solution's installation stack.

3. Choose the **Outputs** tab. The domain name appears as the value for the **ApiEndpoint** key. This URL is the endpoint URL for your newly provisioned image handler API.

To use the solution with your frontend application, use the following example syntax for the API call:

```
https://<ApiEndpoint>/<encodedRequest>
```

Create and use image requests

This solution generates a CloudFront domain name that gives you access to both original and modified images through the image handler API. You can specify parameters such as the image's location and edits to be made in a JSON object on the frontend.

Follow these step-by-step instructions to create image requests:

1. Retrieve your API endpoint for the solution. Refer to Use the solution with a frontend application for instructions.

2. In a code sandbox, or in your frontend application, create a new `imageRequest` JSON object. This object contains the key-value pairs needed to successfully retrieve and perform edits on your images. Using the following code sample and the sharp documentation, adjust the following properties to meet your image editing requirements.

   - **Bucket** – Specify the S3 bucket containing your original image file. This is the name that's specified in the **SourceBuckets** template parameter. You can update the image location by adding it into the SOURCE_BUCKETS environment variable of your image handler Lambda function. See Using AWS Lambda environment variables in the *AWS Lambda Developer Guide* for more information.

   - **Key** – Specify the filename of your original image. This name should include the file extension and subfolders between its location and the root of the bucket. For example, `folder1/folder2/image.jpeg`.

   - **Edits** – Specify image edits as key-value pairs. If you don't specify image edits, the original image returns with no changes made.

   For example, the following code block specifies the image location as `myImageBucket` and specifies edits of `grayscale: true` to change the image to grayscale:

```
const imageRequest = JSON.stringify({
    bucket: "<myImageBucket>",
    key: "<myImage.jpeg>",
    edits: {
        grayscale: true
    }
})
```

3. Stringify the JSON request object. For example:

```
const stringifiedObject = JSON.stringify(<myObject>);
```

4. Base64 encode the JSON string. For example:

```
const encodedObject = btoa(<stringifiedObject>);
```

5. Append the encoded string onto the CloudFront URL. For example:

```
const url = '${<ApiEndpoint>}/${<encodedObject>}';
```

6. Use that URL either in the JavaScript as part of a GET request, or in the frontend as part of an HTML img tag's src property.

For information regarding how to use additional features in an image request, refer to Use smart cropping, Use round cropping, and Activate and use content moderation. For additional features supported by sharp, refer to the sharp documentation.

> **ⓘ Note**
>
> The following filters are not supported for multi-page GIF images due to limitations in the underlying libraries: **rotate**, **smartCrop**, **roundCrop**, and **contentModeration**.

Dynamically resize photos

This solution offers the following **fit** options to dynamically resize an image: cover, contain, fill, inside, and outside. Refer to the sharp documentation for a description of each fit. For example:

```
const imageRequest = JSON.stringify({
    bucket: "<myImageBucket>",
    key: "<myImage.jpeg>",
    edits: {
        resize: {
            width: 200,
            height: 250,
            fit: "cover"
        }
    }
})
```

If you use contain as the resize **fit** mode, you can specify the color of the fill by providing the hex code of the color you want to use. For example:

```
const imageRequest = JSON.stringify({
    bucket: "<myImageBucket>",
    key: "<myImage.jpeg>",
```

```
    edits: {
        resize: {
            width: 200,
            height: 250,
            fit: "contain",
            background: {
                r: 255,
                g: 0,
                b: 0,
                alpha: 1
            }
        }
    }
})
```

Edit images

You can use this solution to edit your images, such as rotating them or changing the coloring to negative. Refer to the [sharp documentation](#) for a description of each operation. For example, to produce a negative of an image, enter the following:

```
const imageRequest = JSON.stringify({
    bucket: "<myImageBucket>",
    key: "<myImage.jpeg>",
    edits: {
        negate: true
    }
})
```

Use smart cropping

This solution uses Amazon Rekognition for face detection in images submitted for smart cropping. To activate smart cropping on an image, add the **smartCrop** property to the **edits** property in the [image request](#).

- **smartCrop (optional, boolean || object)** – Activates the smart cropping feature for an original image. If the value is `true`, then the feature returns the first face detected from the original image with no additional options. For example:

```
const imageRequest = JSON.stringify({
    bucket: "<myImageBucket>",
    key: "<myImage.jpeg>",
```

```
        edits: {
            smartCrop: true
        }
})
```

The following **smartCrop** variables are shown in the following code sample:

- **smartCrop.faceIndex (optional, number)** – Specifies which face to focus on if multiple are present within an original image. The solution indexes detected faces in a zero-based array from the largest detected face to the smallest. If this value isn't specified, Amazon Rekognition returns the largest face detected from the original image.

- **smartCrop.padding (optional, number)** – Specifies an amount of padding in pixels to add around the cropped image. The solution applies the padding value to all sides of the cropped image.

```
const imageRequest = JSON.stringify({
    bucket: "<myImageBucket>",
    key: "<myImage.jpeg>",
    edits: {
        smartCrop: {
            faceIndex: 1,   // zero-based index of detected faces
            padding: 40,    // padding expressed in pixels, applied to all sides
        }
    }
})
```

> ⓘ **Note**
>
> **smartCrop** is not supported for animated (such as, GIF) images.

Use round cropping

This solution can crop images in a circular pattern. To activate round cropping on an image, add the **roundCrop** property to the **edits** property in the image request.

- **roundCrop (optional, boolean || object)** – Activates the round cropping feature for an original image. If the value is true, then the feature returns a circular cropped image that's centered from the original image and has a diameter of the smallest edge of the original image. For example:

```
const imageRequest = JSON.stringify({
    bucket: "<myImageBucket>",
    key: "<myImage.jpeg>",
    edits: {
        roundCrop: true
    }
})
```

The following **roundCrop** variables are shown in the following code sample:

- **roundCrop.rx (optional, number)** – Specifies the radius along the x-axis of the ellipse. If a value isn't provided, the image handler defaults to a value that's half the length of the smallest edge.

- **roundCrop.ry (optional, number)** – Specifies the radius along the y-axis of the ellipse. If a value isn't provided, the image handler defaults to a value that's half the length of the smallest edge.

- **roundCrop.top (optional, number)** – Specifies the offset from the top of the original image to place the center of the ellipse. If a value isn't provided, the image handler defaults to a value that's half of the height.

- **roundCrop.left (optional, number)** – Specifies the offset from the left-most edge of the original image to place the center of the ellipse. If a value isn't provided, the image handler defaults to a value that's half of the width.

```
const imageRequest = JSON.stringify({
    bucket: "<myImageBucket>",
    key: "<myImage.jpeg>",
    edits: {
        roundCrop: {
            rx: 30,   // x-axis radius
            ry: 20,   // y-axis radius
            top: 300, // offset from top edge of original image
            left: 500 // offset from left edge of original image
        }
    }
})
```

> **ⓘ Note**
>
> **roundCrop** is not supported for animated (such as, GIF) images.

Overwrite animated stubs

This solution assumes that GIF files with multiple pages should be animated. If you'd like to indicate that a GIF should not be animated, or that another file type should be animated, include the animated property in the edits property in the image request.

- **animated (optional, boolean)** – Overwrites the initial animated status of the image. If the value is `true` , the solution will attempt to process the image as animated. For example:

```
const imageRequest = JSON.stringify({
    bucket: "<myImageBucket>",
    key: "<myImage.webp>",
    edits: {
        animated: true
    }
})
```

If it is `false`, the solution will process the image as a still image. For example:

```
const imageRequest = JSON.stringify({
    bucket: "<myImageBucket>",
    key: "<myImage.gif>",
    edits: {
        animated: false
    }
})
```

> **ⓘ Note**
>
> If an image does not have multiple pages, it will always be processed as still, regardless of the **edits.animated** property. The following filters are not supported for images that are animated: **rotate**, **smartCrop**, **roundCrop**, and **contentModeration**.

Activate and use content moderation

This solution can detect inappropriate content using Amazon Rekognition. To activate content moderation, add the **contentModeration** property to the **edits** property in the image request.

- **contentModeration (optional, boolean || object)** – Activates the content moderation feature for an original image. If the value is true, then the feature detects inappropriate content using Amazon Rekognition with a minimum confidence that's set higher than 75%. If Amazon Rekognition finds inappropriate content, the solution blurs the image. For example:

```
const imageRequest = JSON.stringify({
    bucket: "<myImageBucket>",
    key: "<myImage.jpeg>",
    edits: {
        contentModeration: true
    }
})
```

The following **contentModeration** variables are shown in the following code sample:

- **contentModeration.minConfidence (optional, number)** – Specifies the minimum confidence level for Amazon Rekognition to use. Amazon Rekognition only returns detected content that's higher than the minimum confidence. If a value isn't provided, the default value is set to 75%.

- **contentModeration.blur (optional, number)** – Specifies the intensity level that an image is blurred if inappropriate content is found. The number represents the sigma of the Gaussian mask, where *sigma = 1 + radius /2*. For more information, refer to the [sharp](#) documentation. If a value isn't provided, the default value is set to 50.

- **contentModeration.moderationLabels (optional, array)** – Identifies the specific content to search for. The image is blurred only if Amazon Rekognition locates the content specified in the **smartCrop.moderationLabels** provided. You can use either a top-level category or a second-level category. Top-level categories include its associated second-level categories. For more information about moderation label options, refer to [Content moderation](#) in the *Amazon Rekognition Developer Guide*.

```
const imageRequest = JSON.stringify({
    bucket: "<myImageBucket>",
    key: "<myImage.jpeg>",
    edits: {
        contentModeration: {
```

```
              minConfidence: 90,   // minimum confidence level for inappropriate
   content

              blur: 80,              // amount to blur image
              moderationLabels: [ // labels to search for
                 "Hate Symbols",
                 "Smoking"
               ]
         }
      }
})
```

> ℹ️ **Note**
>
> **contentModeration** is not supported for animated (such as, GIF) images.

Use supported Thumbor filters

This solution supports the Thumbor filters listed in this section, using API calls. To retrieve your API endpoint for the solution, refer [Use the solution with a frontend application](#) for instructions.

To use the filters, use the following example syntax for the API call:

```
https://<ApiEndpoint>/<modification>/<image.jpeg>
```

Define the source bucket for the request

To define the bucket used when getting the image for a request, include **s3:BucketName** as a modification in your request. For example, if your source buckets were "test-bucket-1, the-other-test-bucket", to indicate that the-other-test-bucket should be used when processing an image, enter the following:

```
https://<ApiEndpoint>/s3:the-other-test-bucket/<image.jpeg>
```

> ℹ️ **Note**
>
> Using the **s3:BucketName** tag requires that the bucket chosen is part of the **SourceBuckets** provided upon deployment. For information on how to change the **SourceBuckets** after deployment, see [the section called "Backward compatibility"](#).

Resize an image

To resize an image, specify `fit-in` and the desired image size. For example, to resize a JPEG image to 300 pixels wide and 400 pixels tall, enter the following:

```
https://<ApiEndpoint>/fit-in/<300x400>/<image.jpeg>
```

Use filters

To use filters, specify a filter from the following table. For example, to blur a JPEG image, enter the following:

```
https://<ApiEndpoint>/filters:blur(7)/<image.jpeg>
```

> ℹ️ **Note**
>
> Some Thumbor filters aren't supported in the current version of this solution. This might affect legacy users with advanced image request configurations. For notes about Thumbor compatibility and source image storage limitations, see [the section called "Backward compatibility"](#). For examples of filter usage, refer to the [Thumbor documentation](#).

Filter name	Filter syntax
Animated	`/filters:animated(true/false)/`
Autojpg	`/filters:autojpg()/`
Background color	`/filters:background_color(color)/`

Filter name	Filter syntax
Blur	`/filters:blur(7)/`
Color fill	`/filters:fill(color)/`
Convolution	`/filters:convolution(1;2;1; 2;4;2;1;2;1,3,false)/`
Crop	`/10x10:100x100/`
Equalize	`/filters:equalize()/`
Grayscale	`/filters:grayscale()/`
Image format (.gif, .heic, .heif, .jpeg, .png, .av	`/filters:format(image_format)`
No upscale	`/filters:no_upscale()/`
Proportion	`/filters:proportion(0.0-1.0)/` `,`
Quality	`/filters:quality(0-100)/`
Resize	`/fit-in/800x1000/`
RGB	`/filters:rgb(20,-20,40)/`
Rotate	`/filters:rotate(90)/`
Sharpen	`/filters:sharpen(0.0-10.0, 0.0-2.0, true/false)/`
Stretch	`/filters:stretch()/`
Strip Exif	`/filters:strip_exif()/`
Strip ICC	`/filters:strip_icc()/`
Upscale	`/filters:upscale()/`

Filter name	Filter syntax

Watermark	`/filters:watermark(bucket,k`
	`ey,x,y,alpha[,w_ratio[,h_ra`
	`tio]])`

Use multiple filters

To use multiple filters on an image, list them in the same section of the URL. Filters process the image in the order that you specify them. For example:

```
https://<api-endpoint>/fit-in/<300x400>/filters:<fill>(<00ff00>)/
filters:<rotate>(<90>)/<image.jpeg>
```

Custom image requests

You can customize most settings for this solution by editing and updating the environment variables associated with the image handler Lambda function. You can find the image handler function in the AWS Management Console using one of the following methods:

Using the AWS Lambda console:

1. Sign in to the AWS Lambda console.
2. Select **Functions**. The image handler function is listed with the following naming convention: `<StackName>-ImageHandlerFunction-<UniqueID>`.

Using the AWS CloudFormation console:

1. Sign in to the AWS CloudFormation console.
2. On the **Stacks** page, select this solution's installation stack.
3. Choose the **Resources** tab. The image handler function is listed with a **Logical ID** of `ImageHandlerFunction`.

After opening the Lambda function, go to the **Environment variables** section. Use the following key-value pairs to customize the solutions settings.

> **ⓘ Note**
>
> The solution uses the [template parameter inputs](#) to determine these initial key values, except for **REWRITE_MATCH_PATTERN** and **REWRITE_SUBSTITUTION**.

Variable Key	Value Type	Description
AUTO_WEPB	`Yes/No`	Choose whether to automatically accept webp image formats.
CORS_ENABLED	`Yes/No`	Indicates whether to return an **Access-Control-Allow-Origin** header with the image handler API response.
CORS_ORIGIN	`String`	This value is returned by the API in the **Access-Control-Allow-Origin** header. An asterisk (*) value supports any origin. We recommend specifying a specific origin (for example, `http://`*`<example>`* `.`*`<domain>`*) to restrict cross-site access to your API. > **ⓘ Note** > This value is ignored if **CORS_ENABLED** is set to No.
ENABLE_DEFAULT_FALLBACK_IMAGE	`Yes/No`	Choose whether to return the default fallback image when errors occur.

Variable Key	Value Type	Description
DEFAULT_FALLBACK_IMAGE_BUCKET	String	Specifies the S3 bucket which contains the default fallback image. ⓘ **Note** This value is ignored if the **ENABLE_DEFAULT_FALLBACK_IMAGE** parameter is set to No.
DEFAULT_FALLBACK_IMAGE_KEY	String	Defines the default fallback image S3 object key, including the prefix. ⓘ **Note** This value is ignored if the **ENABLE_DEFAULT_FALLBACK_IMAGE** parameter is set to No.
ENABLE_SIGNATURE	Yes/No	Choose whether to use the image URL signature.

Variable Key	Value Type	Description
REWRITE_MATCH_PATTERN	`Regex`	By default, this parameter is empty. If you overwrite this default value, use a JavaScript-compatible regular expression for matching custom image requests using the rewrite function. This value should match the JavaScript compatible regular expression. For example, `/(filters-)/gm` .
REWRITE_SUBSTITUTION	`String`	By default, this parameter is empty. If you overwrite this default value, use a substitution string for custom image requests using the rewrite function. For example, `filters:` .
SECRETS_MANAGER	`String`	Defines the Secrets Manager secret that contains the secret key for the image URL signature. ⓘ **Note** This value is ignored if *ENABLE_SIGNATURE* is set to No.

Variable Key	Value Type	Description
SECRET_KEY	`String`	Defines the Secrets Manager secret key that contains the secret value to create the image URL signature. ⓘ **Note** **Note:** This value is ignored if **ENABLE_SI GNATURE** is set to No.
SOURCE_BUCKETS	`String`	The S3 bucket (or buckets) in your account that contain(s) the original images. If you're providing multiple buckets, separate them by commas.

Use the rewrite feature

You can use this solution's rewrite feature to migrate your current image request model to the Serverless Image Handler solution, without changing the applications to accommodate new image URLs.

The rewrite feature translates custom URL image requests into Thumbor-consumable formats, based on JavaScript-compatible regular expression match patterns and substitution strings. After the image request is converted into Thumbor-consumable form, it's then processed as a Thumbor image request and edits are mapped to the new sharp image library.

This feature requires that you populate the following environment variables in the image handler function. These environment variables are added to the function by default, but are left empty for user input if the rewrite feature is needed.

Variable Key	Value Type	Description
REWRITE_MATCH_PATTERN	`Regex`	By default, this parameter is empty. If you overwrite this default value, use a JavaScript-compatible regular expression for matching custom image requests using the rewrite function. This value should match the JavaScript compatible regular expression. For example, `/(filters-)/gm` .
REWRITE_SUBSTITUTION	`String`	By default, this parameter is empty. If you overwrite this default value, use a substitution string for custom image requests using the rewrite function. For example, `filters:`.

You can use any of the Thumbor-supported filters listed in this section with the rewrite feature. The following sections provide examples.

Replace filters- with filters:

If you put `/(filters-)/gm` in REWRITE_MATCH_PATTERN and `filters:` in **REWRITE_SUBSTITUTION**, you can call

```
https://<your-CloudFront-distribution>/filters:rotate(90)/<your-image>
```

instead of

```
https://<your-CloudFront-distribution>/filters-rotate(90)/<your-image>
```

to rotate your image. In this example, the solution replaces `filters-` (filters hyphen syntax) with `filters:` (filters colon syntax).

Reverse path order

You can place filters at the end of the path rather than before the image key.

1. Use the REWRITE_MATCH_PATTERN with a regular expression that parses the path into two groups. The solution then uses REWRITE_SUBSTITUTION to switch the order of the groups.

2. Use a regular expression specified by REWRITE_MATCH_PATTERN to parse the path into groups for a request like `https://abcd.cloudfront.net/imagekey.png/fit-in/200x200`, where the image key appears before the filters. For example:

```
REWRITE_MATCH_PATTERN = /^\/(.*?\..*?)\/(.+)$/gm
```

3. Reverse the order of the fields with REWRITE_SUBSTITUTION to convert the request into a Thumbor style request like `https://abcd.cloudfront.net/fit-in/200x200/imagekey.png`, where the image key is moved to the end of the request. For example:

```
REWRITE_SUBSTITUTION = /$2/$1
```

Parse request type

Refer to [image-request.spec.js.file](#), in the Serverless Image Handler GitHub repository.

Rotate images manually

Images containing rotational EXIF data might not be rotated if the image isn't a JPEG file. Not all browsers support rotational EXIF data for all image formats. If the images aren't a JPEG file, you might need to modify the solution to manually rotate the image based on the EXIF data. To modify the solutions you can access orientation using the [sharp input metadata](#). You can use the orientation data to [rotate the image accordingly](#).

Developer guide

This section provides the source code and an API reference for the solution.

Source code

Visit our GitHub repository to download the source files for this solution and to share your customizations with others. Additionally, if you require an earlier version of the CloudFormation template, you can request from the GitHub issues page.

The AWS CDK generates the Serverless Image Handler templates. See the README.md file for additional information.

API reference

This uses the `sharp` Node.js library to provide high-speed image processing. Open the library, then select **API** from the navigation menu to view the API guides.

Reference

This section includes information about an optional feature for collecting unique metrics for this solution, pointers to related resources, and a list of builders who contributed to this solution.

Anonymized data collection

This solution includes an option to send anonymized operational metrics to AWS. We use this data to better understand how customers use this solution and related services and products. When invoked, the following information is collected and sent to AWS:

- **Solution ID** - The AWS solution identifier
- **Version** – The Serverless Image Handler solution version
- **Unique ID (UUID)** – Randomly generated, unique identifier
- **Timestamp** – The timestamp when the solution's Lambda function runs
- **Region** – The AWS Region the solution is being deployed in
- **CorsEnable** – Whether CORS is activated
- **NumberOfSourceBuckets** – Number of source buckets
- **DeployDemoUi** – Whether the Demo UI deployment is activated
- **LogRetentionPeriod** – The log retention period
- **AutoWebP** – Whether AutoWebP is activated
- **EnableSignature** – Whether the image URL signature is activated
- **EnableDefaultFallbackImage** – Whether the default fallback image is activated

AWS owns the data gathered though this survey. Data collection is subject to the AWS Privacy Notice. To opt out of this feature, complete the following steps before launching the CloudFormation template.

1. Download the `serverless-image-handler.template` the section called "AWS CloudFormation template" to your local hard drive.
2. Open the CloudFormation template with a text editor.
3. Modify the CloudFormation template for the resource with the name starting with `CommonResourcesCustomResourcesCustomResourceAnonymousMetric*` from:

```
AnonymousData: "Yes"
```

to:

```
AnonymousData: "No"
```

4. Sign in to the AWS CloudFormation console.

5. Select **Create stack**.

6. On the **Create stack** page, **Specify template** section, select **Upload a template file**.

7. Under **Upload a template file**, choose **Choose file** and select the edited template from your local drive.

8. Choose **Next** and follow the steps in Deployment process overview to launch the solution.

Related resources

Refer to the sharp Node.js library for more information about sharp.

Contributors

- Simon Krol
- Kamyar Ziabari
- Ryan Hayes
- Beomseok Lee
- George Lenz
- Dmitry Fisenko
- Doug Toppin
- Garvit Singh

Revisions

Date	Change
June 2017	Initial release
August 2017	Solution updated to add the rewrite feature and the optional deployment of a demo UI
October 2017	Solution updated to provide CORS support
September 2018	Added information on watermarking, URL encoding, debugging, and troubleshooting
December 2018	Added information about the Amazon CloudFront distribution for the static website hosted in the Amazon S3 bucket
January 2019	Added information about using the demo UI, safe URLs, and customizing the Thumbor Lambda package
June 2019	Added support for sharp, multiple image sources, basic image editing, smartcrop ping with Amazon Rekognition, backward compatibility, and refresh of demo UI
August 2019	Updated the list of supported Thumbor filters
December 2019	Added information on support for Node.js update
February 2020	Added watermark support for Thumbor filter; added AutoWebP parameter for viewing webp image formats automatically.

Date	Change
August 2020	Updated the AWS CloudFormation template; for more information, refer to the CHANGELOG.md file in the GitHub repository.
November 2020	Added the image URL signature and the default fallback image features; for more information, refer to the CHANGELOG.md file in the GitHub repository.
February 2023	Release v6.1.1: Added package-lock.json, github workflows, demo-ui unicode support, package updates for axios, json5, contribut ing guidelines. For more information on new features, refer to the CHANGELOG.md file in the GitHub repository.
April 2023	Release v6.1.2: s3 bucket ownership control permission and object ownership to custom resource for CloudFront logging bucket, changed xml2js version to 0.5.0. For more information on new features, refer to the CHANGELOG.md file in the GitHub repository.
October 2023	Release v6.2.3: Security patch. For more information on new features, refer to the CHANGELOG.md file in the GitHub repository.
January 2021	Release v5.2.0: Added content moderation, round crop, and support for opt-in Regions; for more information on new features, refer to the CHANGELOG.md file in the GitHub repository
May 2021	Document enhancements to provide clearer business value and better describe architect ure flow

Date	Change
December 2021	Release v6.0.0: Used AWS Cloud Development Kit (AWS CDK) to create the AWS CloudFormation template; added CROP feature in Thumbor URLs; for more information on new features, refer to the [CHANGELOG.md](#) file in the GitHub repository
November 2022	Release v6.1.0: Added Service Catalog AppRegistry, TIFF, and GIF support. Upgraded to AWS CDK v2 and Node 16 Lambda runtime support. For more information on new features, refer to the [CHANGELOG.md](#) file in the GitHub repository
February 2023	Release v6.1.1: Added `package-lock.json` , GitHub workflows, demo-ui unicode support, package updates for axios, json5, and contributing guidelines. For more information on updates and new features, refer to the [CHANGELOG.md](#) file in the GitHub repository.
April 2023	Release v6.1.2: Mitigated impact caused by new default settings for S3 Object Ownership (ACLs disabled) for all new S3 buckets. For more information on new features, refer to the [CHANGELOG.md](#) file in the GitHub repository.
July 2023	Updated documentation template to improve readability and navigation. Created a **Use the solution** section for a user's guide. Provided more detailed examples for image editing and Thumbor filters.

Date	Change
August 2023	Release v6.2.0: Add `cdk-helper` module, `cdk-deploy` enhancements, and bug fixes.For more information on updates and new features, refer to the [CHANGELOG.md](#) file in the GitHub repository.
August 2023	Release v6.2.1: Addressed deployment issue with demo UI disabled, and other bug fixes. For more information on updates and new features, refer to the [CHANGELOG.md](#) file in the GitHub repository.
September 2023	Release v6.2.2: Upgraded Sharp to v0.32.6 for vulnerability CVE-2023-4863. Upgraded outdated NPM packages, and cleaned up some licensing and documentation. For more information on updates and new features, refer to the [CHANGELOG.md](#) file in the GitHub repository.
October 2023	Release v6.2.3: Security patch. For more information on updates and new features, refer to the [CHANGELOG.md](#) file in the GitHub repository.
November 2023	Documentation update: Added [Confirm cost tags associated with the solution](#) to the Monitoring the solution with AWS Service Catalog AppRegistry section.
December 2023	Release v6.2.4: Updated CDK to 2.111.0. Updated Node 20.x Lambda runtimes. Enabled Thumbor filter chaining. For more informati on on updates and new features, refer to the [CHANGELOG.md](#) file in the GitHub repository.

Date	Change
December 2023	Documentation update: Clarified [rewrite feature](#) instructions.
January 2024	Release 6.2.5: Addressed image metadata when generating Amazon Rekognition compatible images. Exclude `demo-ui-config` from being deleted when updating to a new version. Changed error return for invalid overlay image size. GIF files without multiple pages are now treated as non-animated, allowing all filters to be used on them. Update AWS CDK to 2.118.0 and typescript to 5.3.3. For more information on updates and new features, refer to the [CHANGELOG.md](#) file in the GitHub repository.
June 2024	Release 6.2.6: Added CloudFront logging bucket name to CloudFormation stack outputs; ability to specify buckets for Thumbor-style requests; override for default animated image logic; and type inference for 8-bit depth AVIF images. Decreased permissions allotted to Lambda functions. Dependency updates. Modified JPG type inference to only check first 2 bytes. Fixed Thumbor-style regex matching being overly permissive. Fixed solution version and ID not being passed to backend Lambda function. For more information on updates and new features, refer to the [CHANGELOG.md](#) file in the GitHub repository.

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents AWS current product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers, or licensors. AWS products or services are provided "as is" without warranties, representations, or conditions of any kind, whether express or implied. AWS responsibilities and liabilities to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Serverless Image Handler is licensed under the terms of the Apache License Version 2.0.