亚马逊云科技

User Guide

Amazon ECR

API Version 2015-09-21

Amazon ECR: User Guide

Table of Contents

What is Amazon ECR **1**
Components of Amazon ECR 1
Features of Amazon ECR 2
How to get started with Amazon ECR 3
Pricing for Amazon ECR 3
China regions for Amazon ECR 3
Moving an image through its lifecycle **4**
Prerequisites 4
Install the Amazon CLI 4
Install Docker 4
Step 1: Create a Docker image 6
Step 2: Authenticate to your default registry 8
Step 3: Create a repository 9
Step 4: Push an image to Amazon ECR 9
Step 5: Pull an image from Amazon ECR 10
Step 6: Delete an image 11
Step 7: Delete a repository 11
Optimizing performance **12**
Private registry **14**
Registry concepts 14
Registry authentication 14
Using the Amazon ECR credential helper 15
Using an authorization token 15
Using HTTP API authentication 16
Registry settings 17
Registry permissions 18
Registry policy examples 18
Granting permissions for cross account replication 21
Granting permissions for pull through cache 23
Private repositories **24**
Repository concepts 24
Creating a repository to store images 25
Next steps 26
Viewing repository details 26

Deleting a repository ... 27
Repository policies ... 28
Repository policies vs IAM policies ... 28
Repository policy examples ... 30
Setting a repository policy statement ... 35
Tagging a repository ... 37
Tag basics ... 37
Tagging your resources for billing ... 37
Adding tags ... 38
Deleting tags ... 39
Private images ... 41
Pushing an image ... 41
Required IAM permissions ... 42
Pushing a Docker image ... 43
Pushing a multi-architecture image ... 45
Pushing a Helm chart ... 46
Signing an image ... 49
Considerations ... 49
Prerequisites ... 49
Configure authentication for the Notary client ... 50
Signing an image ... 50
Next steps ... 51
Deleting a signature ... 51
Viewing image details ... 52
Pulling an image ... 53
Pulling the Amazon Linux container image ... 54
Deleting an image ... 55
Retagging an image ... 56
Preventing image tags from being overwritten ... 59
Setting image tag mutability (Amazon Web Services Management Console) ... 59
Setting image tag mutability (Amazon CLI) ... 60
Container image manifest formats ... 60
Amazon ECR image manifest conversion ... 61
Using Amazon ECR images with Amazon ECS ... 62
Required IAM permissions ... 62
Specifying an Amazon ECR image in a task definition ... 64

Using Amazon ECR Images with Amazon EKS ... 64
Required IAM permissions ... 65
Installing a Helm chart on an Amazon EKS cluster ... 66
Scan images for vulnerabilities ... 68
Filters for repositories ... 69
Filter wildcards ... 69
Enhanced scanning ... 70
Considerations for enhanced scanning ... 70
Required IAM permissions ... 71
Configuring enhanced scanning ... 72
Changing the enhanced scanning duration ... 74
EventBridge events ... 75
Retrieving findings ... 80
Basic scanning ... 81
Region support for improved basic scanning ... 82
Operating system support for basic scanning and improved basic scanning ... 83
Configuring improved basic scanning ... 85
Configuring basic scanning ... 85
Manually scanning an image ... 86
Retrieving findings ... 87
Troubleshooting image scanning ... 88
Understanding scan status `SCAN_ELIGIBILITY_EXPIRED` ... 89
Sync an upstream registry ... 91
Repository creation templates ... 91
Considerations for using pull through cache rules ... 92
Required IAM permissions ... 94
Using registry permissions ... 94
Next steps ... 96
Creating a pull through cache rule ... 96
Prerequisites ... 97
Using the Amazon Web Services Management Console ... 97
Using the Amazon CLI ... 103
Next steps ... 105
Repository creation templates ... 106
How it works ... 106
Required IAM permissions ... 109

Creating a repository creation template ... 110
Deleting a repository creation template ... 112
Validating pull through cache rule ... 112
Pulling an image with a pull through cache rule ... 113
Storing your upstream repository credentials ... 115
Troubleshooting pull through cache issues ... 122
Replicate images ... 124
Considerations for private image replication ... 124
Replication examples ... 125
Example: Configuring cross-Region replication to a single destination Region ... 126
Example: Configuring cross-Region replication using a repository filter ... 126
Example: Configuring cross-Region replication to multiple destination Regions ... 127
Example: Configuring cross-account replication ... 127
Example: Specifying multiple rules in a configuration ... 128
Configuring replication ... 129
Automate the cleanup of images ... 131
How lifecycle policies work ... 131
Lifecycle policy evaluation rules ... 132
Creating a lifecycle policy preview ... 133
Creating a lifecycle policy ... 135
Prerequisite ... 135
Examples of lifecycle policies ... 136
Lifecycle policy template ... 137
Filtering on image age ... 137
Filtering on image count ... 138
Filtering on multiple rules ... 138
Filtering on multiple tags in a single rule ... 141
Filtering on all images ... 143
Lifecycle policy properties ... 146
Rule priority ... 146
Description ... 146
Tag status ... 147
Tag pattern list ... 147
Tag prefix list ... 147
Count type ... 148
Count unit ... 148

Count number ... 149
Action ... 149
Security ... 150
Identity and Access Management ... 150
Audience ... 151
Authenticating with identities ... 152
Managing access using policies ... 154
How Amazon Elastic Container Registry works with IAM ... 156
Identity-based policy examples ... 161
Using Tag-Based Access Control ... 166
Amazon managed policies for Amazon ECR ... 168
Using service-linked roles ... 175
Troubleshooting ... 181
Data protection ... 183
Encryption at rest ... 184
Compliance validation ... 191
Infrastructure Security ... 192
Interface VPC Endpoints (Amazon PrivateLink) ... 193
Cross-service confused deputy prevention ... 201
Monitoring ... 204
Visualizing Your Service Quotas and Setting Alarms ... 205
Usage Metrics ... 206
Usage Reports ... 207
Repository metrics ... 208
Enabling CloudWatch metrics ... 208
Available metrics and dimensions ... 208
Viewing metrics with CloudWatch ... 209
Events and EventBridge ... 209
Sample events from Amazon ECR ... 210
Logging Actions with Amazon CloudTrail ... 214
Amazon ECR information in CloudTrail ... 214
Understanding Amazon ECR log file entries ... 215
Service quotas ... 227
Troubleshooting ... 233
Troubleshooting Docker ... 233
Docker logs do not contain expected error messages ... 233

Error: "Filesystem Verification Failed" or "404: Image Not Found" when pulling an image from an Amazon ECR repository 234
Error: "Filesystem Layer Verification Failed" when pulling images from Amazon ECR 234
HTTP 403 Errors or "no basic auth credentials" error when pushing to repository 235
Troubleshooting Amazon ECR error messages 236
HTTP 429: Too Many Requests or ThrottleException 236
HTTP 403: "User [arn] is not authorized to perform [operation]" 237
HTTP 404: "Repository Does Not Exist" error 237
Error: Cannot perform an interactive login from a non TTY device 237
Document history 239

What is Amazon Elastic Container Registry?

Amazon Elastic Container Registry (Amazon ECR) is an Amazon managed container image registry service that is secure, scalable, and reliable. Amazon ECR supports private repositories with resource-based permissions using Amazon IAM. This is so that specified users or Amazon EC2 instances can access your container repositories and images. You can use your preferred CLI to push, pull, and manage Docker images, Open Container Initiative (OCI) images, and OCI compatible artifacts.

> **Note**
>
> Amazon ECR supports public container image repositories as well. For more information, see What is Amazon ECR Public in the *Amazon ECR Public User Guide*.

The Amazon container services team maintains a public roadmap on GitHub. It contains information about what the teams are working on and allows all Amazon customers the ability to give direct feedback. For more information, see Amazon Containers Roadmap.

Components of Amazon ECR

Amazon ECR contains the following components:

Registry

An Amazon ECR private registry is provided to each Amazon account; you can create one or more repositories in your registry and store Docker images, Open Container Initiative (OCI) images, and OCI compatible artifacts in them. For more information, see Amazon ECR private registry.

Authorization token

Your client must authenticate to an Amazon ECR private registry as an Amazon user before it can push and pull images. For more information, see Private registry authentication in Amazon ECR.

Repository

An Amazon ECR repository contains your Docker images, Open Container Initiative (OCI) images, and OCI compatible artifacts. For more information, see Amazon ECR private repositories.

Repository policy

You can control access to your repositories and the contents within them with repository policies. For more information, see Private repository policies in Amazon ECR.

Image

You can push and pull container images to your repositories. You can use these images locally on your development system, or you can use them in Amazon ECS task definitions and Amazon EKS pod specifications. For more information, see Using Amazon ECR images with Amazon ECS and Using Amazon ECR Images with Amazon EKS.

Features of Amazon ECR

Amazon ECR provides the following features:

- Lifecycle policies help with managing the lifecycle of the images in your repositories. You define rules that result in the cleaning up of unused images. You can test rules before applying them to your repository. For more information, see Automate the cleanup of images by using lifecycle policies in Amazon ECR.
- Image scanning helps in identifying software vulnerabilities in your container images. Each repository can be configured to **scan on push**. This ensures that each new image pushed to the repository is scanned. You can then retrieve the results of the image scan. For more information, see Scan images for software vulnerabilities in Amazon ECR.
- Cross-Region and cross-account replication makes it easier for you to have your images where you need them. This is configured as a registry setting and is on a per-Region basis. For more information, see Private registry settings in Amazon ECR.
- Pull through cache rules provide a way to cache repositories in an upstream registry in your private Amazon ECR registry. Using a pull through cache rule, Amazon ECR will periodically reach out to the upstream registry to ensure the cached image in your Amazon ECR private registry is up to date. For more information, see Sync an upstream registry with an Amazon ECR private registry.

How to get started with Amazon ECR

If you are using Amazon Elastic Container Service (Amazon ECS) or Amazon Elastic Kubernetes Service (Amazon EKS), note that the setup for those two services is similar to the setup for Amazon ECR because Amazon ECR is an extension of both services.

When using the Amazon Command Line Interface with Amazon ECR, use a version of the Amazon CLI that supports the latest Amazon ECR features. If you don't see support for an Amazon ECR feature in the Amazon CLI, upgrade to the latest version of the Amazon CLI. For information about installing the latest version of the Amazon CLI, see Install or update to the latest version of the Amazon CLI in the *Amazon Command Line Interface User Guide*.

To learn how to push a container image to a private Amazon ECR repository using the Amazon CLI and Docker, see Moving an image through its lifecycle in Amazon ECR.

Pricing for Amazon ECR

With Amazon ECR, you only pay for the amount of data you store in your repositories and for the data transfer from your image pushes and pulls. For more information, see Amazon ECR pricing.

China regions for Amazon ECR

Examples throughout this guide reflect regions in different parts of the world. For China, replace those regions with one of the two China regions: cn-northwest-1 or cn-north-1. Also, because there are only two regions in China, any issues related to three or more regions don't apply to China.

Moving an image through its lifecycle in Amazon ECR

If you are using Amazon ECR for the first time, use the following steps with the Docker CLI and the Amazon CLI to create a sample image, authenticate to the default registry, and create a private repository. Then push an image to and pull an image from the private repository. When you are finished with the sample image, delete the sample image and the repository.

To use the Amazon Web Services Management Console instead of the Amazon CLI, see the section called "Creating a repository to store images".

For more information on the other tools available for managing your Amazon resources, including the different Amazon SDKs, IDE toolkits, and the Windows PowerShell command line tools, see http://www.amazonaws.cn/tools/.

Prerequisites

If you do not have the latest Amazon CLI and Docker installed and ready to use, use the following steps to install both of these tools.

Install the Amazon CLI

To use the Amazon CLI with Amazon ECR, install the latest Amazon CLI version. For information, see Installing the Amazon Command Line Interface in the *Amazon Command Line Interface User Guide*.

Install Docker

Docker is available on many different operating systems, including most modern Linux distributions, like Ubuntu, and even macOS and Windows. For more information about how to install Docker on your particular operating system, go to the Docker installation guide.

You do not need a local development system to use Docker. If you are using Amazon EC2 already, you can launch an Amazon Linux 2023 instance and install Docker to get started.

If you already have Docker installed, skip to Step 1: Create a Docker image.

To install Docker on an Amazon EC2 instance using an Amazon Linux 2023 AMI

1. Launch an instance with the latest Amazon Linux 2023 AMI. For more information, see Launching an instance in the *Amazon EC2 User Guide for Linux Instances*.

2. Connect to your instance. For more information, see Connect to Your Linux Instance in the *Amazon EC2 User Guide for Linux Instances*.
3. Update the installed packages and package cache on your instance.

```
sudo yum update -y
```

4. Install the most recent Docker Community Edition package.

```
sudo yum install docker
```

5. Start the Docker service.

```
sudo service docker start
```

6. Add the `ec2-user` to the `docker` group so you can execute Docker commands without using `sudo`.

```
sudo usermod -a -G docker ec2-user
```

7. Log out and log back in again to pick up the new `docker` group permissions. You can accomplish this by closing your current SSH terminal window and reconnecting to your instance in a new one. Your new SSH session will have the appropriate `docker` group permissions.
8. Verify that the `ec2-user` can run Docker commands without `sudo`.

```
docker info
```

> **Note**
>
> In some cases, you may need to reboot your instance to provide permissions for the `ec2-user` to access the Docker daemon. Try rebooting your instance if you see the following error:
>
> ```
> Cannot connect to the Docker daemon. Is the docker daemon running on this
> host?
> ```

Step 1: Create a Docker image

In this step, you create a Docker image of a simple web application, and test it on your local system or Amazon EC2 instance.

To create a Docker image of a simple web application

1. Create a file called `Dockerfile`. A Dockerfile is a manifest that describes the base image to use for your Docker image and what you want installed and running on it. For more information about Dockerfiles, go to the Dockerfile Reference.

```
touch Dockerfile
```

2. Edit the `Dockerfile` you just created and add the following content.

```
FROM public.ecr.aws/amazonlinux/amazonlinux:latest

# Install dependencies
RUN yum update -y && \
 yum install -y httpd

# Install apache and write hello world message
RUN echo 'Hello World!' > /var/www/html/index.html

# Configure apache
RUN echo 'mkdir -p /var/run/httpd' >> /root/run_apache.sh && \
 echo 'mkdir -p /var/lock/httpd' >> /root/run_apache.sh && \
 echo '/usr/sbin/httpd -D FOREGROUND' >> /root/run_apache.sh && \
 chmod 755 /root/run_apache.sh

EXPOSE 80

CMD /root/run_apache.sh
```

   This Dockerfile uses the public Amazon Linux 2 image hosted on Amazon ECR Public. The `RUN` instructions update the package caches, installs some software packages for the web server, and then write the "Hello World!" content to the web servers document root. The `EXPOSE` instruction exposes port 80 on the container, and the `CMD` instruction starts the web server.

3. Build the Docker image from your Dockerfile.

> **Note**
>
> Some versions of Docker may require the full path to your Dockerfile in the following command, instead of the relative path shown below.

```
docker build -t hello-world .
```

4. List your container image.

```
docker images --filter reference=hello-world
```

Output:

```
REPOSITORY          TAG                 IMAGE ID            CREATED
 SIZE
hello-world         latest              e9ffedc8c286        4 minutes ago
 194MB
```

5. Run the newly built image. The `-p 80:80` option maps the exposed port 80 on the container to port 80 on the host system. For more information about **docker run**, go to the Docker run reference.

```
docker run -t -i -p 80:80 hello-world
```

> **Note**
>
> Output from the Apache web server is displayed in the terminal window. You can ignore the "`Could not reliably determine the fully qualified domain name`" message.

6. Open a browser and point to the server that is running Docker and hosting your container.

   - If you are using an EC2 instance, this is the **Public DNS** value for the server, which is the same address you use to connect to the instance with SSH. Make sure that the security group for your instance allows inbound traffic on port 80.
   - If you are running Docker locally, point your browser to http://localhost/.

- If you are using **docker-machine** on a Windows or Mac computer, find the IP address of the VirtualBox VM that is hosting Docker with the **docker-machine ip** command, substituting *machine-name* with the name of the docker machine you are using.

```
docker-machine ip machine-name
```

You should see a web page with your "Hello World!" statement.

7. Stop the Docker container by typing **Ctrl + c**.

Step 2: Authenticate to your default registry

After you have installed and configured the Amazon CLI, authenticate the Docker CLI to your default registry. That way, the **docker** command can push and pull images with Amazon ECR. The Amazon CLI provides a **get-login-password** command to simplify the authentication process.

To authenticate Docker to an Amazon ECR registry with get-login-password, run the **aws ecr get-login-password** command. When passing the authentication token to the **docker login** command, use the value `AWS` for the username and specify the Amazon ECR registry URI you want to authenticate to. If authenticating to multiple registries, you must repeat the command for each registry.

> **Important**
>
> If you receive an error, install or upgrade to the latest version of the Amazon CLI. For more information, see Installing the Amazon Command Line Interface in the *Amazon Command Line Interface User Guide*.

- get-login-password (Amazon CLI)

```
aws ecr get-login-password --region region | docker login --username AWS --password-
stdin aws_account_id.dkr.ecr.region.amazonaws.com
```

- Get-ECRLoginCommand (Amazon Tools for Windows PowerShell)

```
(Get-ECRLoginCommand).Password | docker login --username AWS --password-
stdin aws_account_id.dkr.ecr.region.amazonaws.com
```

Step 3: Create a repository

Now that you have an image to push to Amazon ECR, you must create a repository to hold it. In this example, you create a repository called `hello-repository` to which you later push the `hello-world:latest` image. To create a repository, run the following command:

```
aws ecr create-repository \
    --repository-name hello-repository \
    --region region
```

Step 4: Push an image to Amazon ECR

Now you can push your image to the Amazon ECR repository you created in the previous section. Use the **docker** CLI to push images after the following prerequisites are met:

- The minimum version of **docker** is installed: 1.7.
- The Amazon ECR authorization token was configured with **docker login**.
- The Amazon ECR repository exists and the user has access to push to the repository.

After those prerequisites are met, you can push your image to your newly created repository in the default registry for your account.

To tag and push an image to Amazon ECR

1. List the images you have stored locally to identify the image to tag and push.

   ```
   docker images
   ```

   Output:

   ```
   REPOSITORY          TAG                 IMAGE ID            CREATED
    VIRTUAL SIZE
   hello-world         latest              e9ffedc8c286        4 minutes ago
    241MB
   ```

2. Tag the image to push to your repository.

```
docker tag hello-world:latest aws_account_id.dkr.ecr.region.amazonaws.com/hello-
repository
```

3. Push the image.

```
docker push aws_account_id.dkr.ecr.region.amazonaws.com/hello-repository
```

Output:

```
The push refers to a repository [aws_account_id.dkr.ecr.region.amazonaws.com/hello-
repository] (len: 1)
e9ae3c220b23: Pushed
a6785352b25c: Pushed
0998bf8fb9e9: Pushed
0a85502c06c9: Pushed
latest: digest: sha256:215d7e4121b30157d8839e81c4e0912606fca105775bb0636EXAMPLE
 size: 6774
```

Step 5: Pull an image from Amazon ECR

After your image is pushed to your Amazon ECR repository, you can pull it from other locations. Use the **docker** CLI to pull images after the following prerequisites are met:

- The minimum version of **docker** is installed: 1.7.
- The Amazon ECR authorization token was configured with **docker login**.
- The Amazon ECR repository exists and the user has access to pull from the repository.

After those prerequisites are met, you can pull your image. To pull your example image from Amazon ECR, run the following command:

```
docker pull aws_account_id.dkr.ecr.region.amazonaws.com/hello-repository:latest
```

Output:

```
latest: Pulling from hello-repository
0a85502c06c9: Pull complete
0998bf8fb9e9: Pull complete
```

```
a6785352b25c: Pull complete
e9ae3c220b23: Pull complete
Digest: sha256:215d7e4121b30157d8839e81c4e0912606fca105775bb0636EXAMPLE
Status: Downloaded newer image for aws_account_id.dkr.region.amazonaws.com/hello-
repository:latest
```

Step 6: Delete an image

If you no longer need an image in one of your repositories, you can delete the image. To delete an image, specify the repository that it's in and either an `imageTag` or `imageDigest` value for the image. The following example deletes an image in the `hello-repository` repository with the image tag `latest`. To delete your example image from the repository, run the following command:

```
aws ecr batch-delete-image \
      --repository-name hello-repository \
      --image-ids imageTag=latest \
      --region region
```

Step 7: Delete a repository

If you no longer need an entire repository of images, you can delete the repository. The following example uses the `--force` flag to delete a repository that contains images. To delete a repository that contains images (and all the images within it), run the following command:

```
aws ecr delete-repository \
      --repository-name hello-repository \
      --force \
      --region region
```

Optimizing performance for Amazon ECR

You can use the following recommendations about settings and strategies to optimize performance when using Amazon ECR.

Use Docker 1.10 and above to take advantage of simultaneous layer uploads

Docker images are composed of layers, which are intermediate build stages of the image. Each line in a Dockerfile results in the creation of a new layer. When you use Docker 1.10 and above, Docker defaults to pushing as many layers as possible as simultaneous uploads to Amazon ECR, resulting in faster upload times.

Use a smaller base image

The default images available through Docker Hub may contain many dependencies that your application doesn't require. Consider using a smaller image created and maintained by others in the Docker community, or build your own base image using Docker's minimal scratch image. For more information, see Create a base image in the Docker documentation.

Place the dependencies that change the least earlier in your Dockerfile

Docker caches layers, and that speeds up build times. If nothing on a layer has changed since the last build, Docker uses the cached version instead of rebuilding the layer. However, each layer is dependent on the layers that came before it. If a layer changes, Docker recompiles not only that layer, but any layers that come after that layer as well.

To minimize the time required to rebuild a Dockerfile and to re-upload layers, consider placing the dependencies that change the least frequently earlier in your Dockerfile. Place rapidly changing dependencies (such as your application's source code) later in the stack.

Chain commands to avoid unnecessary file storage

Intermediate files created on a layer remain a part of that layer even if they are deleted in a subsequent layer. Consider the following example:

```
WORKDIR /tmp
RUN wget http://example.com/software.tar.gz
RUN wget tar -xvf software.tar.gz
RUN mv software/binary /opt/bin/myapp
RUN rm software.tar.gz
```

In this example, the layers created by the first and second RUN commands contain the original .tar.gz file and all of its unzipped contents. This is even though the .tar.gz file is deleted by the fourth RUN command. These commands can be chained together into a single RUN statement to ensure that these unnecessary files aren't part of the final Docker image:

```
WORKDIR /tmp
RUN wget http://example.com/software.tar.gz &&\
    wget tar -xvf software.tar.gz &&\
    mv software/binary /opt/bin/myapp &&\
    rm software.tar.gz
```

Use the closest regional endpoint

You can reduce latency in pulling images from Amazon ECR by ensuring that you are using the regional endpoint closest to where your application is running. If your application is running on an Amazon EC2 instance, you can use the following shell code to obtain the region from the Availability Zone of the instance:

```
REGION=$(curl -s http://169.254.169.254/latest/meta-data/placement/availability-zone
 |\
     sed -n 's/\(\d*\)[a-zA-Z]*$/\1/p')
```

The region can be passed to Amazon CLI commands using the **--region** parameter, or set as the default region for a profile using the **aws configure** command. You can also set the region when making calls using the Amazon SDK. For more information, see the documentation for the SDK for your specific programming language.

Amazon ECR private registry

An Amazon ECR private registry hosts your container images in a highly available and scalable architecture. You can use your private registry to manage private image repositories consisting of Docker and Open Container Initiative (OCI) images and artifacts. Each Amazon account is provided with a default private Amazon ECR registry. For more information about Amazon ECR public registries, see Public registries in the *Amazon Elastic Container Registry Public User Guide*.

Private registry concepts

- The URL for your default private registry is `https://aws_account_id.dkr.ecr.cn-north-1.amazonaws.com.cn`.
- By default, your account has read and write access to the repositories in your private registry. However, users require permissions to make calls to the Amazon ECR APIs and to push or pull images to and from your private repositories. Amazon ECR provides several managed policies to control user access at varying levels. For more information, see Amazon Elastic Container Registry Identity-based policy examples.
- You must authenticate your Docker client to your private registry so that you can use the **docker push** and **docker pull** commands to push and pull images to and from the repositories in that registry. For more information, see Private registry authentication in Amazon ECR.
- Private repositories can be controlled with both user access policies and repository policies. For more information about repository policies, see Private repository policies in Amazon ECR.
- The repositories in your private registry can be replicated across Regions in your own private registry and across separate accounts by configuring replication for your private registry. For more information, see Private image replication in Amazon ECR.

Private registry authentication in Amazon ECR

You can use the Amazon Web Services Management Console, the Amazon CLI, or the Amazon SDKs to create and manage private repositories. You can also use those methods to perform some actions on images, such as listing or deleting them. These clients use standard Amazon authentication methods. Even though you can use the Amazon ECR API to push and pull images, you're more likely to use the Docker CLI or a language-specific Docker library.

The Docker CLI doesn't support native IAM authentication methods. Additional steps must be taken so that Amazon ECR can authenticate and authorize Docker push and pull requests.

The registry authentication methods that are detailed in the following sections are available.

Using the Amazon ECR credential helper

Amazon ECR provides a Docker credential helper which makes it easier to store and use Docker credentials when pushing and pulling images to Amazon ECR. For installation and configuration steps, see Amazon ECR Docker Credential Helper.

> **Note**
>
> The Amazon ECR Docker credential helper doesn't support multi-factor authentication (MFA) currently.

Using an authorization token

An authorization token's permission scope matches that of the IAM principal used to retrieve the authentication token. An authentication token is used to access any Amazon ECR registry that your IAM principal has access to and is valid for 12 hours. To obtain an authorization token, you must use the GetAuthorizationToken API operation to retrieve a base64-encoded authorization token containing the username `AWS` and an encoded password. The Amazon CLI `get-login-password` command simplifies this by retrieving and decoding the authorization token which you can then pipe into a **docker login** command to authenticate.

To authenticate Docker to an Amazon ECR private registry with get-login

- To authenticate Docker to an Amazon ECR registry with get-login-password, run the **aws ecr get-login-password** command. When passing the authentication token to the **docker login** command, use the value `AWS` for the username and specify the Amazon ECR registry URI you want to authenticate to. If authenticating to multiple registries, you must repeat the command for each registry.

> **Important**
>
> If you receive an error, install or upgrade to the latest version of the Amazon CLI. For more information, see Installing the Amazon Command Line Interface in the *Amazon Command Line Interface User Guide*.

- get-login-password (Amazon CLI)

```
aws ecr get-login-password --region region | docker login --username AWS --
password-stdin aws_account_id.dkr.ecr.region.amazonaws.com
```

- Get-ECRLoginCommand (Amazon Tools for Windows PowerShell)

```
(Get-ECRLoginCommand).Password | docker login --username AWS --password-
stdin aws_account_id.dkr.ecr.region.amazonaws.com
```

Using HTTP API authentication

Amazon ECR supports the Docker Registry HTTP API. However, because Amazon ECR is a private registry, you must provide an authorization token with every HTTP request. You can add an HTTP authorization header using the `-H` option for **curl** and pass the authorization token provided by the **get-authorization-token** Amazon CLI command.

To authenticate with the Amazon ECR HTTP API

1. Retrieve an authorization token with the Amazon CLI and set it to an environment variable.

```
TOKEN=$(aws ecr get-authorization-token --output text --query
 'authorizationData[].authorizationToken')
```

2. To authenticate to the API, pass the `$TOKEN` variable to the `-H` option of **curl**. For example, the following command lists the image tags in an Amazon ECR repository. For more information, see the Docker Registry HTTP API reference documentation.

```
curl -i -H "Authorization: Basic $TOKEN"
 https://aws_account_id.dkr.ecr.region.amazonaws.com/v2/amazonlinux/tags/list
```

The output is as follows:

```
HTTP/1.1 200 OK
Content-Type: text/plain; charset=utf-8
Date: Thu, 04 Jan 2018 16:06:59 GMT
Docker-Distribution-Api-Version: registry/2.0
Content-Length: 50
Connection: keep-alive

{"name":"amazonlinux","tags":["2017.09","latest"]}
```

Private registry settings in Amazon ECR

Amazon ECR uses private registry settings to configure features at the registry level. The private registry settings are configured separately for each Region. You can use private registry settings to configure the following features.

- **Registry permissions**—A registry permissions policy provides control over the replication and pull through cache permissions. For more information, see Private registry permissions in Amazon ECR.
- **Pull through cache rules**—A pull through cache rule is used to cache images from an upstream registry in your Amazon ECR private registry. For more information, see Sync an upstream registry with an Amazon ECR private registry.
- **Replication configuration**—The replication configuration is used to control whether your repositories are copied across Amazon Regions or accounts. For more information, see Private image replication in Amazon ECR
- **Repository creation templates**—A repository creation template is used to define the standard settings to apply when new repositories are created by Amazon ECR on your behalf. For example, repositories created by a pull through cache action. For more information, see Templates to control repositories created during a pull through cache action.
- **Scanning configuration**—By default, your registry is enabled for basic scanning. You may enable enhanced scanning which provides an automated, continuous scanning mode that scans for both operating system and programming language package vulnerabilities. For more information, see Scan images for software vulnerabilities in Amazon ECR.

Private registry permissions in Amazon ECR

Amazon ECR uses a **registry policy** to grant permissions to an Amazon principal at the private registry level. These permissions are used to scope access to the replication and pull through cache features.

Amazon ECR only enforces the following permissions at the private registry level. If any additional actions are added to the registry policy, an error will occur.

- `ecr:ReplicateImage` – Grants permission to another account, referred to as the source registry, to replicate its images to your registry. This is only used for cross-account replication.
- `ecr:BatchImportUpstreamImage` – Grants permission to retrieve the external image and import it to your private registry.
- `ecr:CreateRepository` – Grants permission to create a repository in a private registry. This permission is required if the repository storing either the replicated or cached images doesn't already exist in the private registry.

> **Note**
>
> While it is possible to add the `ecr:*` action to a private registry permissions policy, it is considered best practice to only add the specific actions required based on the feature you're using rather than use a wildcard.

Topics

- Private registry policy examples for Amazon ECR
- Granting registry permissions for cross account replication in Amazon ECR
- Granting registry permissions for pull through cache in Amazon ECR

Private registry policy examples for Amazon ECR

The following examples show registry permissions policy statements that you could use to control the permissions that users have to your Amazon ECR registry.

> **Note**
>
> In each example, if the `ecr:CreateRepository` action is removed from your registry permission statement, replication can still occur. However, for successful replication, you need to create repositories with the same name within your account.

Example: Allow the root user of a source account to replicate all repositories

The following registry permissions policy allows the root user of a source account to replicate all repositories.

```
{
    "Version":"2012-10-17",
    "Statement":[
        {
            "Sid":"ReplicationAccessCrossAccount",
            "Effect":"Allow",
            "Principal":{
                "AWS":"arn:aws:iam::source_account_id:root"
            },
            "Action":[
                "ecr:CreateRepository",
                "ecr:ReplicateImage"
            ],
            "Resource": [
                "arn:aws:ecr:us-west-2:your_account_id:repository/*"
            ]
        }
    ]
}
```

Example: Allow root users from multiple accounts

The following registry permissions policy has two statements. Each statement allows the root user of a source account to replicate all repositories.

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

```
            "Sid":"ReplicationAccessCrossAccount",
            "Effect":"Allow",
            "Principal":{
                "AWS":"arn:aws:iam::source_account_id:root"
            },
            "Action":[
                "ecr:CreateRepository",
                "ecr:ReplicateImage"
            ],
            "Resource": [
                "arn:aws:ecr:us-west-2:your_account_id:repository/*"
            ]
        },
        {
            "Sid":"ReplicationAccessCrossAccount",
            "Effect":"Allow",
            "Principal":{
                "AWS":"arn:aws:iam::source_account_id:root"
            },
            "Action":[
                "ecr:CreateRepository",
                "ecr:ReplicateImage"
            ],
            "Resource": [
                "arn:aws:ecr:us-west-2:your_account_id:repository/*"
            ]
        }
    ]
}
```

Example: Allow the root user of a source account to replicate all repositories with prefix `prod-`.

The following registry permissions policy allows the root user of a source account to replicate all repositories that start with `prod-`.

```
{
    "Version":"2012-10-17",
    "Statement":[
        {
            "Sid":"ReplicationAccessCrossAccount",
            "Effect":"Allow",
            "Principal":{
```

```
                "AWS":"arn:aws:iam::source_account_id:root"
            },
            "Action":[
                "ecr:CreateRepository",
                "ecr:ReplicateImage"
            ],
            "Resource": [
                "arn:aws:ecr:us-west-2:your_account_id:repository/prod-*"
            ]
        }
    ]
}
```

Granting registry permissions for cross account replication in Amazon ECR

The cross account policy type is used to grant permissions to an Amazon principal, allowing the replication of the repositories from a source registry to your registry. By default, you have permission to configure cross-Region replication within your own registry. You only need to configure the registry policy if you're granting another account permission to replicate contents to your registry.

A registry policy must grant permission for the `ecr:ReplicateImage` API action. This API is an internal Amazon ECR API that can replicate images between Regions or accounts. You can also grant permission for the `ecr:CreateRepository` permission, which allows Amazon ECR to create repositories in your registry if they don't exist already. If the `ecr:CreateRepository` permission isn't provided, a repository with the same name as the source repository must be created manually in your registry. If neither is done, replication fails. Any failed CreateRepository or ReplicateImage API actions show up in CloudTrail.

To configure a permissions policy for replication (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.
2. From the navigation bar, choose the Region to configure your registry policy in.
3. In the navigation pane, choose **Private registry**, **Registry permissions**.
4. On the **Registry permissions** page, choose **Generate statement**.
5. Complete the following steps to define your policy statement using the policy generator.

    a. For **Policy type**, choose **Cross account policy**.

b. For **Statement ID**, enter a unique statement ID. This field is used as the `Sid` on the registry policy.

c. For **Accounts**, enter the account IDs for each account you want to grant permissions to. When specifying multiple account IDs, separate them with a comma.

6. Expand the **Preview policy statement** section to review the registry permissions policy statement.

7. After the policy statement is confirmed, choose **Add to policy** to save the policy to your registry.

To configure a permissions policy for replication (Amazon CLI)

1. Create a file named `registry_policy.json` and populate it with a registry policy.

```
{
    "Version":"2012-10-17",
    "Statement":[
        {
            "Sid":"ReplicationAccessCrossAccount",
            "Effect":"Allow",
            "Principal":{
                "AWS":"arn:aws:iam::source_account_id:root"
            },
            "Action":[
                "ecr:CreateRepository",
                "ecr:ReplicateImage"
            ],
            "Resource": [
                "arn:aws:ecr:us-west-2:your_account_id:repository/*"
            ]
        }
    ]
}
```

2. Create the registry policy using the policy file.

```
aws ecr put-registry-policy \
      --policy-text file://registry_policy.json \
      --region us-west-2
```

3. Retrieve the policy for your registry to confirm.

```
aws ecr get-registry-policy \
      --region us-west-2
```

Granting registry permissions for pull through cache in Amazon ECR

Amazon ECR private registry permissions may be used to scope the permissions of individual IAM entities to use pull through cache. If an IAM entity has more permissions granted by an IAM policy than the registry permissions policy is granting, the IAM policy takes precedence.

To create a private registry permissions policy (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.
2. From the navigation bar, choose the Region to configure your private registry permissions statement in.
3. In the navigation pane, choose **Private registry**, **Registry permissions**.
4. On the **Registry permissions** page, choose **Generate statement**.
5. For each pull through cache permissions policy statement you want to create, do the following.

   a. For **Policy type**, choose **Pull through cache policy**.

   b. For **Statement id**, provide a name for the pull through cache statement policy.

   c. For **IAM entities**, specify the users, groups, or roles to include in the policy.

   d. For **Repository namespace**, select the pull through cache rule to associate the policy with.

   e. For **Repository names**, specify the repository base name to apply the rule for. For example, if you want to specify the Amazon Linux repository on Amazon ECR Public, the repository name would be `amazonlinux`.

Amazon ECR private repositories

An Amazon ECR private repository contains your Docker images, Open Container Initiative (OCI) images, and OCI compatible artifacts. You can create, monitor, and delete image repositories and set permissions that control who can access them by using Amazon ECR API operations or the **Repositories** section of the Amazon ECR console. Amazon ECR also integrates with the Docker CLI, so that you can push and pull images from your development environments to your repositories.

Topics

- Private repository concepts
- Creating an Amazon ECR private repository to store images
- Viewing the contents and details of a private repository in Amazon ECR
- Deleting a private repository in Amazon ECR
- Private repository policies in Amazon ECR
- Tagging a private repository in Amazon ECR

Private repository concepts

- By default, your account has read and write access to the repositories in your default registry (`aws_account_id.dkr.ecr.region.amazonaws.com`). However, users require permissions to make calls to the Amazon ECR APIs and to push or pull images to and from your repositories. Amazon ECR provides several managed policies to control user access at varying levels. For more information, see Amazon Elastic Container Registry Identity-based policy examples.
- Repositories can be controlled with both user access policies and individual repository policies. For more information, see Private repository policies in Amazon ECR.
- Repository names can support namespaces, which you can use to group similar repositories. For example, if there are several teams using the same registry, Team A can use the `team-a` namespace, and Team B can use the `team-b` namespace. By doing this, each team has their own image called `web-app` with each image prefaced with the team namespace. This configuration allows these images on each team to be used simultaneously without interference. Team A's image is `team-a/web-app`, and Team B's image is `team-b/web-app`.
- Your images can be replicated to other repositories across Regions in your own registry and across accounts. You can do this by specifying a replication configuration in your registry settings. For more information, see Private registry settings in Amazon ECR.

Creating an Amazon ECR private repository to store images

Create an Amazon ECR private repository, and then use the repository to store your container images. Use the following steps to create a private repository using the Amazon Web Services Management Console. For steps to create a repository using the Amazon CLI, see Step 3: Create a repository.

To create a repository (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region to create your repository in.
3. On the **Repositories** page, choose **Private repositories**, and then choose **Create repository**.
4. For **Visibility settings**, verify that **Private** is selected.
5. For **Repository name**, enter a unique name for your repository. The repository name can be specified on its own (for example `nginx-web-app`). Alternatively, it can be prepended with a namespace to group the repository into a category (for example `project-a/nginx-web-app`).

   > **Note**
   >
   > The repository name may container a maximum of 256 characters. The name must start with a letter and can only contain lowercase letters, numbers, hyphens, underscores, periods and forward slashes. Using a double hyphen, double underscore, or double forward slash isn't supported.

6. For **Tag immutability**, choose the tag mutability setting for the repository. Repositories configured with immutable tags prevent image tags from being overwritten. For more information, see Preventing image tags from being overwritten in Amazon ECR.
7. For **Scan on push**, while you can specify the scan settings at the repository level for basic scanning, it is best practice to specify the scan configuration at the private registry level. Specify the scanning settings at the private registry allow you to enable either enhanced scanning or basic scanning as well as define filters to specify which repositories are scanned. For more information, see Scan images for software vulnerabilities in Amazon ECR.
8. For **KMS encryption**, choose whether to enable encryption of the images in the repository using Amazon Key Management Service. By default, when KMS encryption is enabled, Amazon ECR uses an Amazon managed key (KMS key) with the alias `aws/ecr`. This key is created in

your account the first time that you create a repository with KMS encryption enabled. For more information, see Encryption at rest.

9. When KMS encryption is enabled, select **Customer encryption settings (advanced)** to choose your own KMS key. The KMS key must be in the same Region as the cluster. Choose **Create an Amazon KMS key** to navigate to the Amazon KMS console to create your own key.
10. Choose **Create repository**.

Next steps

To view the steps to push an image to your repository, select the repository and choose **View push commands**. For more information about pushing an image to your repository, see Pushing an image to an Amazon ECR private repository.

Viewing the contents and details of a private repository in Amazon ECR

After you created a private repository, you can view details about the repository in the Amazon Web Services Management Console:

- Which images are stored in a repository
- Details about each image stored in the repository, including the size and SHA digest for each image
- The scan frequency specified for the contents of the repository
- Whether the repository has an active pull through cache rule associated with it
- The encryption setting for the repository

> **Note**
>
> Starting with Docker version 1.9, the Docker client compresses image layers before pushing them to a V2 Docker registry. The output of the **docker images** command shows the uncompressed image size. Therefore, keep in mind that Docker might return a larger image than the image shown in the Amazon Web Services Management Console.

To view repository information (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region that contains the repository to view.
3. In the navigation pane, choose **Repositories**.
4. On the **Repositories** page, choose the **Private** tab and then the repository to view.
5. On the repository detail page, the console defaults to the **Images** view. Use the navigation menu to view other information about the repository.
   - Choose **Summary** to view the repository details and pull count data for the repository.
   - Choose **Images** to view information about the image tags in the repository. To view more information about the image, select the image tag. For more information, see Viewing image details in Amazon ECR.

     If there are untagged images that you want to delete, you can select the box to the left of the repositories to delete and choose **Delete**. For more information, see Deleting an image in Amazon ECR.
   - Choose **Permissions** to view the repository policies that are applied to the repository. For more information, see Private repository policies in Amazon ECR.
   - Choose **Lifecycle Policy** to view the lifecycle policy rules that are applied to the repository. The lifecycle events history is also viewed here. For more information, see Automate the cleanup of images by using lifecycle policies in Amazon ECR.
   - Choose **Tags** to view the metadata tags that are applied to the repository.

Deleting a private repository in Amazon ECR

If you're finished using a repository, you can delete it. When you delete a repository in the Amazon Web Services Management Console, all of the images contained in the repository are also deleted; this cannot be undone.

> **Important**
>
> Images in the deleted repositories are also deleted. You cannot undo this operation.

To delete a repository (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region that contains the repository to delete.
3. In the navigation pane, choose **Repositories**.
4. On the **Repositories** page, choose the **Private** tab and then select the repository to delete and choose **Delete**.
5. In the **Delete** ***repository_name*** window, verify that the selected repositories should be deleted and choose **Delete**.

Private repository policies in Amazon ECR

Amazon ECR uses resource-based permissions to control access to repositories. Resource-based permissions let you specify which users or roles have access to a repository and what actions they can perform on the repository. By default, only the Amazon account that created the repository has access to the repository. You can apply a repository policy that allows additional access to your repository.

Topics

- Repository policies vs IAM policies
- Private repository policy examples in Amazon ECR
- Setting a private repository policy statement in Amazon ECR

Repository policies vs IAM policies

Amazon ECR repository policies are a subset of IAM policies that are scoped for, and specifically used for, controlling access to individual Amazon ECR repositories. IAM policies are generally used to apply permissions for the entire Amazon ECR service but can also be used to control access to specific resources as well.

Both Amazon ECR repository policies and IAM policies are used when determining which actions a specific user or role may perform on a repository. If a user or role is allowed to perform an action through a repository policy but is denied permission through an IAM policy (or vice versa) then the action will be denied. A user or role only needs to be allowed permission for an action through either a repository policy or an IAM policy but not both for the action to be allowed.

> **Important**
>
> Amazon ECR requires that users have permission to make calls to the `ecr:GetAuthorizationToken` API through an IAM policy before they can authenticate to a registry and push or pull any images from any Amazon ECR repository. Amazon ECR provides several managed IAM policies to control user access at varying levels; for more information, see Amazon Elastic Container Registry Identity-based policy examples.

You can use either of these policy types to control access to your repositories, as shown in the following examples.

This example shows an Amazon ECR repository policy, which allows for a specific user to describe the repository and the images within the repository.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ECRRepositoryPolicy",
            "Effect": "Allow",
            "Principal": {"AWS": "arn:aws:iam::account-id:user/username"},
            "Action": [
                "ecr:DescribeImages",
                "ecr:DescribeRepositories"
            ]
        }
    ]
}
```

This example shows an IAM policy that achieves the same goal as above, by scoping the policy to a repository (specified by the full ARN of the repository) using the resource parameter. For more information about Amazon Resource Name (ARN) format, see Resources.

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

```
            "Action": [
                "ecr:DescribeImages",
                "ecr:DescribeRepositories"
            ],
            "Resource": ["arn:aws:ecr:region:account-id:repository/repository-name"]
        }
    ]
}
```

Private repository policy examples in Amazon ECR

> **Important**
>
> The repository policy examples on this page are meant to be applied to Amazon ECR private repositories. They will not work properly if used with an IAM principal directly unless modified to specify the Amazon ECR repository as the resource. For more information on setting repository policies, see Setting a private repository policy statement in Amazon ECR.

Amazon ECR repository policies are a subset of IAM policies that are scoped for, and specifically used for, controlling access to individual Amazon ECR repositories. IAM policies are generally used to apply permissions for the entire Amazon ECR service but can also be used to control access to specific resources as well. For more information, see Repository policies vs IAM policies.

The following repository policy examples show permission statements that you could use to control access to your Amazon ECR private repositories.

> **Important**
>
> Amazon ECR requires that users have permission to make calls to the `ecr:GetAuthorizationToken` API through an IAM policy before they can authenticate to a registry and push or pull any images from any Amazon ECR repository. Amazon ECR provides several managed IAM policies to control user access at varying levels; for more information, see Amazon Elastic Container Registry Identity-based policy examples.

Example: Allow one or more users

The following repository policy allows one or more users to push and pull images to and from a repository.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowPushPull",
            "Effect": "Allow",
            "Principal": {
                "AWS": [
                    "arn:aws:iam::account-id:user/push-pull-user-1",
                    "arn:aws:iam::account-id:user/push-pull-user-2"
                ]
            },
            "Action": [
                "ecr:BatchGetImage",
                "ecr:BatchCheckLayerAvailability",
                "ecr:CompleteLayerUpload",
                "ecr:GetDownloadUrlForLayer",
                "ecr:InitiateLayerUpload",
                "ecr:PutImage",
                "ecr:UploadLayerPart"
            ]
        }
    ]
}
```

Example: Allow another account

The following repository policy allows a specific account to push images.

> **Important**
>
> The account you are granting permissions to must have the Region you are creating the repository policy in enabled, otherwise an error will occur.

```
{
    "Version": "2012-10-17",
```

```
    "Statement": [
        {
            "Sid": "AllowCrossAccountPush",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::account-id:root"
            },
            "Action": [
                "ecr:BatchCheckLayerAvailability",
                "ecr:CompleteLayerUpload",
                "ecr:InitiateLayerUpload",
                "ecr:PutImage",
                "ecr:UploadLayerPart"
            ]
        }
    ]
}
```

The following repository policy allows some users to pull images (*pull-user-1* and *pull-user-2*) while providing full access to another (*admin-user*).

> **Note**
>
> For more complicated repository policies that are not currently supported in the Amazon Web Services Management Console, you can apply the policy with the **set-repository-policy** Amazon CLI command.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowPull",
            "Effect": "Allow",
            "Principal": {
                "AWS": [
                    "arn:aws:iam::account-id:user/pull-user-1",
                    "arn:aws:iam::account-id:user/pull-user-2"
                ]
            },
            "Action": [
                "ecr:BatchGetImage",
```

```
                "ecr:GetDownloadUrlForLayer"
            ]
        },
        {
            "Sid": "AllowAll",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::account-id:user/admin-user"
            },
            "Action": [
                "ecr:*"
            ]
        }
    ]
}
```

Example: Deny all

The following repository policy denies all users in all accounts the ability to pull images.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "DenyPull",
            "Effect": "Deny",
            "Principal": "*",
            "Action": [
                "ecr:BatchGetImage",
                "ecr:GetDownloadUrlForLayer"
            ]
        }
    ]
}
```

Example: Restricting access to specific IP addresses

The following example denies permissions to any user to perform any Amazon ECR operations when applied to a repository from a specific range of addresses.

The condition in this statement identifies the `54.240.143.*` range of allowed Internet Protocol version 4 (IPv4) IP addresses.

The `Condition` block uses the `NotIpAddress` conditions and the `aws:SourceIp` condition key, which is an Amazon-wide condition key. For more information about these condition keys, see Amazon Global Condition Context Keys. The `aws:sourceIp` IPv4 values use the standard CIDR notation. For more information, see IP Address Condition Operators in the *IAM User Guide*.

```
{
    "Version": "2012-10-17",
    "Id": "ECRPolicyId1",
    "Statement": [
        {
            "Sid": "IPAllow",
            "Effect": "Deny",
            "Principal": "*",
            "Action": "ecr:*",
            "Condition": {
                "NotIpAddress": {
                    "aws:SourceIp": "54.240.143.0/24"
                }
            }
        }
    ]
}
```

Example: Allow an Amazon service

The following repository policy allows Amazon CodeBuild access to the Amazon ECR API actions necessary for integration with that service. When using the following example, you should use the `aws:SourceArn` and `aws:SourceAccount` condition keys to scope which resources can assume these permissions. For more information, see Amazon ECR sample for CodeBuild in the *Amazon CodeBuild User Guide*.

```
{
   "Version":"2012-10-17",
   "Statement":[
      {
         "Sid":"CodeBuildAccess",
         "Effect":"Allow",
         "Principal":{
            "Service":"codebuild.amazonaws.com"
         },
         "Action":[
            "ecr:BatchGetImage",
```

```
            "ecr:GetDownloadUrlForLayer"
         ],
         "Condition":{
            "ArnLike":{
               "aws:SourceArn":"arn:aws:codebuild:region:123456789012:project/project-
name"
            },
            "StringEquals":{
               "aws:SourceAccount":"123456789012"
            }
         }
      }
   ]
}
```

Setting a private repository policy statement in Amazon ECR

You can add an access policy statement to a repository in the Amazon Web Services Management Console by following the steps below. You can add multiple policy statements per repository. For example policies, see Private repository policy examples in Amazon ECR.

> **Important**
>
> Amazon ECR requires that users have permission to make calls to the `ecr:GetAuthorizationToken` API through an IAM policy before they can authenticate to a registry and push or pull any images from any Amazon ECR repository. Amazon ECR provides several managed IAM policies to control user access at varying levels; for more information, see Amazon Elastic Container Registry Identity-based policy examples.

To set a repository policy statement

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region that contains the repository to set a policy statement on.
3. In the navigation pane, choose **Repositories**.
4. On the **Repositories** page, choose the repository to set a policy statement on to view the contents of the repository.
5. From the repository image list view, in the navigation pane, choose **Permissions**, **Edit**.

> **Note**
>
> If you don't see the **Permissions** option in the navigation pane, ensure that you are in the repository image list view.

6. On the **Edit permissions** page, choose **Add statement**.
7. For **Statement name**, enter a name for the statement.
8. For **Effect**, choose whether the policy statement will result in an allow or an explicit deny.
9. For **Principal**, choose the scope to apply the policy statement to. For more information, see Amazon JSON Policy Elements: Principal in the *IAM User Guide*.

   - You can apply the statement to all authenticated Amazon users by selecting the **Everyone (*)** check box.
   - For **Service principal**, specify the service principal name (for example, `ecs.amazonaws.com`) to apply the statement to a specific service.
   - For **Amazon Account IDs**, specify an Amazon account number (for example, `111122223333`) to apply the statement to all users under a specific Amazon account. Multiple accounts can be specified by using a comma delimited list.

     > **Important**
     >
     > The account you are granting permissions to must have the Region you are creating the repository policy in enabled, otherwise an error will occur.

   - For **IAM Entities**, select the roles or users under your Amazon account to apply the statement to.

     > **Note**
     >
     > For more complicated repository policies that are not currently supported in the Amazon Web Services Management Console, you can apply the policy with the **set-repository-policy** Amazon CLI command.

10. For **Actions**, choose the scope of the Amazon ECR API operations that the policy statement should apply to from the list of individual API operations.
11. When you are finished, choose **Save** to set the policy.

12. Repeat the previous step for each repository policy to add.

Tagging a private repository in Amazon ECR

To help you manage your Amazon ECR repositories, you can assign your own metadata to new or existing Amazon ECR repositories by using Amazon resource *tags*. For example, you could define a set of tags for your account's Amazon ECR repositories that helps you track the owner of each repository.

Tag basics

Tags don't have any semantic meaning to Amazon ECR and are interpreted strictly as a string of characters. Tags are not automatically assigned to your resources. You can edit tag keys and values, and you can remove tags from a resource at any time. You can set the value of a tag to an empty string, but you can't set the value of a tag to null. If you add a tag that has the same key as an existing tag on that resource, the new value overwrites the old value. If you delete a resource, any tags for the resource are also deleted.

You can work with tags using the Amazon ECR console, the Amazon CLI, and the Amazon ECR API.

Using Amazon Identity and Access Management (IAM), you can control which users in your Amazon account have permission to create, edit, or delete tags. For information about tags in IAM policies, see the section called “Using Tag-Based Access Control”.

Tagging your resources for billing

The tags you add to your Amazon ECR repositories are helpful when reviewing cost allocation after enabling them in your Cost & Usage Report. For more information, see Amazon ECR usage reports.

To see the cost of your combined resources, you can organize your billing information based on resources that have the same tag key values. For example, you can tag several resources with a specific application name, and then organize your billing information to see the total cost of that application across several services. For more information about setting up a cost allocation report with tags, see The Monthly Cost Allocation Report in the *Amazon Billing User Guide*.

> **Note**
>
> If you've just enabled reporting, data for the current month is available for viewing after 24 hours.

Adding tags to a private repository in Amazon ECR

You can add tags to a private repository.

For information about names and best practices for tags, see Tag naming limits and requirements and Best practices in the *Tagging Amazon Resources User Guide*.

Adding tags to a repository (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.
2. From the navigation bar, select the region to use.
3. In the navigation pane, choose **Repositories**.
4. On the **Repositories** page, select the check box next to the repository you want to tag.
5. From the **Action** menu, select **Repository tags**.
6. On the **Repository tags** page, select **Add tags**, **Add tag**.
7. On the **Edit repository tags** page, specify the key and value for each tag, and then choose **Save**.

Adding tags to a repository (Amazon CLI or API)

You can add or overwrite one or more tags by using the Amazon CLI or an API.

- Amazon CLI - tag-resource
- API action - TagResource

The following examples show how to add tags using the Amazon CLI.

Example 1: Tag a repository

The following command tags a repository.

```
aws ecr tag-resource \
     --resource-arn arn:aws:ecr:region:account_id:repository/repository_name \
     --tags Key=stack,Value=dev
```

Example 2: Tag a repository with multiple tags

The following command adds three tags to a repository.

```
aws ecr tag-resource \
     --resource-arn arn:aws:ecr:region:account_id:repository/repository_name \
     --tags Key=key1,Value=value1 Key=key2,Value=value2 Key=key3,Value=value3
```

Example 3: List tags for a repository

The following command lists the tags associated with a repository.

```
aws ecr list-tags-for-resource \
     --resource-arn arn:aws:ecr:region:account_id:repository/repository_name
```

Example 4: Create a repository and add a tag

The following command creates a repository named `test-repo` and adds a tag with key `team` and value `devs`.

```
aws ecr create-repository \
     --repository-name test-repo \
     --tags Key=team,Value=devs
```

Deleting tags from a private repository in Amazon ECR

You can delete tags from a private repository.

To delete a tag from a private repository (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.
2. From the navigation bar, select the region to use.
3. On the **Repositories** page, select the check box next to the repository you want to remove a tag from.
4. From the **Action** menu, select **Repository tags**.
5. On the **Repository tags** page, select **Edit**.
6. On the **Edit repository tags** page, select **Remove** for each tag you want to delete, and choose **Save**.

To delete a tag from a private repository (Amazon CLI)

You can delete one or more tags by using the Amazon CLI or an API.

- Amazon CLI - untag-resource
- API action - UntagResource

The following example shows how to delete a tag from a repository using the Amazon CLI.

```
aws ecr untag-resource \
     --resource-arn arn:aws:ecr:region:account_id:repository/repository_name \
     --tag-keys tag_key
```

Private images in Amazon ECR

Amazon ECR stores Docker images, Open Container Initiative (OCI) images, and OCI compatible artifacts in private repositories. You can use the Docker CLI, or your preferred client, to push and pull images to and from your repositories.

Topics

- Pushing an image to an Amazon ECR private repository
- Signing an image stored in an Amazon ECR private repository
- Deleting a signature from an Amazon ECR private repository
- Viewing image details in Amazon ECR
- Pulling an image to your local environment from an Amazon ECR private repository
- Pulling the Amazon Linux container image
- Deleting an image in Amazon ECR
- Retagging an imagein Amazon ECR
- Preventing image tags from being overwritten in Amazon ECR
- Container image manifest format support in Amazon ECR
- Using Amazon ECR images with Amazon ECS
- Using Amazon ECR Images with Amazon EKS

Pushing an image to an Amazon ECR private repository

You can push your Docker images, manifest lists, and Open Container Initiative (OCI) images and compatible artifacts to your private repositories.

Amazon ECR also provides a way to replicate your images to other repositories. By specifying a replication configuration in your private registry settings, you can replicate across Regions in your own registry and across different accounts. For more information, see Private registry settings in Amazon ECR.

Topics

- IAM permissions for pushing an image to an Amazon ECR private repository
- Pushing a Docker image to an Amazon ECR private repository

- Pushing a multi-architecture image to an Amazon ECR private repository
- Pushing a Helm chart to an Amazon ECR private repository

IAM permissions for pushing an image to an Amazon ECR private repository

Users need IAM permissions to push images to Amazon ECR private repositories. Following the best practice of granting least privilege, you can grant access to a specific repository. You can also grant access to all repositories.

A user must authenticate to each Amazon ECR registry they want to push images to by requesting an authorization token. Amazon ECR provides several Amazon managed policies to control user access at varying levels. For more information, see Amazon managed policies for Amazon Elastic Container Registry.

You can also create a your own IAM policies. The following IAM policy grants the required permissions for pushing an image to a specific repository. The repository must be specified as a full Amazon Resource Name (ARN).

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "ecr:CompleteLayerUpload",
                "ecr:UploadLayerPart",
                "ecr:InitiateLayerUpload",
                "ecr:BatchCheckLayerAvailability",
                "ecr:PutImage"
            ],
            "Resource": "arn:aws:ecr:region:111122223333:repository/repository-name"
        },
        {
            "Effect": "Allow",
            "Action": "ecr:GetAuthorizationToken",
            "Resource": "*"
        }
    ]
}
```

The following IAM policy grants the required permissions for pushing an image to all repositories.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "ecr:CompleteLayerUpload",
                "ecr:GetAuthorizationToken",
                "ecr:UploadLayerPart",
                "ecr:InitiateLayerUpload",
                "ecr:BatchCheckLayerAvailability",
                "ecr:PutImage"
            ],
            "Resource": "*"
        }
    ]
}
```

Pushing a Docker image to an Amazon ECR private repository

You can push your container images to an Amazon ECR repository with the **docker push** command.

Amazon ECR also supports creating and pushing Docker manifest lists that are used for multi-architecture images. For information, see Pushing a multi-architecture image to an Amazon ECR private repository.

To push a Docker image to an Amazon ECR repository

The Amazon ECR repository must exist before you push the image. For more information, see the section called “Creating a repository to store images”.

1. Authenticate your Docker client to the Amazon ECR registry to which you intend to push your image. Authentication tokens must be obtained for each registry used, and the tokens are valid for 12 hours. For more information, see Private registry authentication in Amazon ECR.

   To authenticate Docker to an Amazon ECR registry, run the **aws ecr get-login-password** command. When passing the authentication token to the **docker login** command, use the value `AWS` for the username and specify the Amazon ECR registry URI you want to authenticate to. If authenticating to multiple registries, you must repeat the command for each registry.

> **Important**
>
> If you receive an error, install or upgrade to the latest version of the Amazon CLI. For more information, see Installing the Amazon Command Line Interface in the *Amazon Command Line Interface User Guide*.

```
aws ecr get-login-password --region region | docker login --username AWS --
password-stdin aws_account_id.dkr.ecr.region.amazonaws.com
```

2. If your image repository doesn't exist in the registry you intend to push to yet, create it. For more information, see Creating an Amazon ECR private repository to store images.
3. Identify the local image to push. Run the **docker images** command to list the container images on your system.

   ```
   docker images
   ```

   You can identify an image with the *repository:tag* value or the image ID in the resulting command output.
4. Tag your image with the Amazon ECR registry, repository, and optional image tag name combination to use. The registry format is *aws_account_id*.dkr.ecr.*cn-north-1*.amazonaws.com.cn. The repository name should match the repository that you created for your image. If you omit the image tag, we assume that the tag is `latest`.

   The following example tags a local image with the ID *e9ae3c220b23* as *aws_account_id*.dkr.ecr.*cn-north-1*.amazonaws.com.cn/my-repository:tag.

   ```
   docker tag e9ae3c220b23 aws_account_id.dkr.ecr.cn-north-1.amazonaws.com.cn/my-
   repository:tag
   ```
5. Push the image using the **docker push** command:

   ```
   docker push aws_account_id.dkr.ecr.cn-north-1.amazonaws.com.cn/my-repository:tag
   ```
6. (Optional) Apply any additional tags to your image and push those tags to Amazon ECR by repeating Step 4 and Step 5.

Pushing a multi-architecture image to an Amazon ECR private repository

You can push multi-architecture images to an Amazon ECR repository by creating and pushing Docker manifest lists. A *manifest list* is a list of images that is created by specifying one or more image names. In most cases, the manifest list is created from images that serve the same function but are for different operating systems or architectures. The manifest list isn't required. For more information, see docker manifest.

A manifest list can be pulled or referenced in an Amazon ECS task definition or Amazon EKS pod spec like other Amazon ECR images.

Prerequisites

- In your Docker CLI, turn on experimental features. For information about experimental features, see Experimental features in the Docker documentation.
- The Amazon ECR repository must exist before you push the image. For more information, see the section called "Creating a repository to store images".
- Images must be pushed to your repository before you create the Docker manifest. For information about how to push an image, see Pushing a Docker image to an Amazon ECR private repository.

To push a multi-architecture Docker image to an Amazon ECR repository

1. Authenticate your Docker client to the Amazon ECR registry to which you intend to push your image. Authentication tokens must be obtained for each registry used, and the tokens are valid for 12 hours. For more information, see Private registry authentication in Amazon ECR.

   To authenticate Docker to an Amazon ECR registry, run the **aws ecr get-login-password** command. When passing the authentication token to the **docker login** command, use the value `AWS` for the username and specify the Amazon ECR registry URI you want to authenticate to. If authenticating to multiple registries, you must repeat the command for each registry.

> **Important**
>
> If you receive an error, install or upgrade to the latest version of the Amazon CLI. For more information, see Installing the Amazon Command Line Interface in the *Amazon Command Line Interface User Guide*.

```
aws ecr get-login-password --region region | docker login --username AWS --
password-stdin aws_account_id.dkr.ecr.region.amazonaws.com
```

2. List the images in your repository, confirming the image tags.

```
aws ecr describe-images --repository-name my-repository
```

3. Create the Docker manifest list. The `manifest create` command verifies that the referenced images are already in your repository and creates the manifest locally.

```
docker manifest create aws_account_id.dkr.ecr.cn-north-1.amazonaws.com.cn/my-
repository aws_account_id.dkr.ecr.cn-north-1.amazonaws.com.cn/my-
repository:image_one_tag aws_account_id.dkr.ecr.cn-north-1.amazonaws.com.cn/my-
repository:image_two
```

4. (Optional) Inspect the Docker manifest list. This enables you to confirm the size and digest for each image manifest referenced in the manifest list.

```
docker manifest inspect aws_account_id.dkr.ecr.cn-north-1.amazonaws.com.cn/my-
repository
```

5. Push the Docker manifest list to your Amazon ECR repository.

```
docker manifest push aws_account_id.dkr.ecr.cn-north-1.amazonaws.com.cn/my-
repository
```

Pushing a Helm chart to an Amazon ECR private repository

You can push Open Container Initiative (OCI) artifacts to an Amazon ECR repository. To see an example of this functionality, use the following steps to push a Helm chart to Amazon ECR.

For information about using your Amazon ECR hosted Helm charts with Amazon EKS, see Installing a Helm chart on an Amazon EKS cluster.

To push a Helm chart to an Amazon ECR repository

1. Install the latest version of the Helm client. These steps were written using Helm version `3.8.2`. For more information, see Installing Helm.
2. Use the following steps to create a test Helm chart. For more information, see Helm Docs - Getting Started.

   a. Create a Helm chart named `helm-test-chart` and clear the contents of the `templates` directory.

   ```
   helm create helm-test-chart
   rm -rf ./helm-test-chart/templates/*
   ```

   b. Create a ConfigMap in the `templates` folder.

   ```
   cd helm-test-chart/templates
   cat <<EOF > configmap.yaml
   apiVersion: v1
   kind: ConfigMap
   metadata:
     name: helm-test-chart-configmap
   data:
     myvalue: "Hello World"
   EOF
   ```

3. Package the chart. The output will contain the filename of the packaged chart which you use when pushing the Helm chart.

   ```
   cd ../..
   helm package helm-test-chart
   ```

   Output

   ```
   Successfully packaged chart and saved it to: /Users/username/helm-test-
   chart-0.1.0.tgz
   ```

4. Create a repository to store your Helm chart. The name of your repository should match the name you used when creating the Helm chart in step 2. For more information, see Creating an Amazon ECR private repository to store images.

```
aws ecr create-repository \
     --repository-name helm-test-chart \
     --region us-west-2
```

5. Authenticate your Helm client to the Amazon ECR registry to which you intend to push your Helm chart. Authentication tokens must be obtained for each registry used, and the tokens are valid for 12 hours. For more information, see Private registry authentication in Amazon ECR.

```
aws ecr get-login-password \
     --region us-west-2 | helm registry login \
     --username AWS \
     --password-stdin aws_account_id.dkr.ecr.cn-north-1.amazonaws.com.cn
```

6. Push the Helm chart using the **helm push** command. The output should include the Amazon ECR repository URI and SHA digest.

```
helm push helm-test-chart-0.1.0.tgz oci://aws_account_id.dkr.ecr.cn-
north-1.amazonaws.com.cn/
```

7. Describe your Helm chart.

```
aws ecr describe-images \
     --repository-name helm-test-chart \
     --region us-west-2
```

In the output, verify that the `artifactMediaType` parameter indicates the proper artifact type.

```
{
    "imageDetails": [
        {
            "registryId": "aws_account_id",
            "repositoryName": "helm-test-chart",
            "imageDigest":
 "sha256:dd8aebdda7df991a0ffe0b3d6c0cf315fd582cd26f9755a347a52adEXAMPLE",
            "imageTags": [
                "0.1.0"
```

```
            ],
            "imageSizeInBytes": 1620,
            "imagePushedAt": "2021-09-23T11:39:30-05:00",
            "imageManifestMediaType": "application/vnd.oci.image.manifest.v1+json",
            "artifactMediaType": "application/vnd.cncf.helm.config.v1+json"
        }
    ]
}
```

8. (Optional) For additional steps, install the Helm configmap and get started with Amazon EKS. For more information, see Installing a Helm chart on an Amazon EKS cluster.

Signing an image stored in an Amazon ECR private repository

Amazon ECR integrates with Amazon Signer to provide a way for you to sign your container images. You can store both your container images and the signatures in your private repositories.

Considerations

The following should be considered when using Amazon ECR image signing.

- Signatures stored in your repository count against the service quota for the maximum number of images per repository. For more information, see Amazon ECR service quotas.
- When using Amazon ECR lifecycle policies, any action by a rule to expire or delete an OCI image index will result in Amazon ECR deleting any signatures referenced by that image index within 24 hours.

Prerequisites

Before you begin, The following prerequisites must be met.

- Install and configure the latest version of the Amazon CLI. For more information, see Installing or updating the latest version of the Amazon CLI in the *Amazon Command Line Interface User Guide*.
- Install the Notation CLI and the Amazon Signer plugin for Notation. For more information, see Prerequisites for signing container images in the *Amazon Signer Developer Guide*.
- Have a container image stored in an Amazon ECR private repository to sign. For more information, see Pushing an image to an Amazon ECR private repository.

Configure authentication for the Notary client

Before you can create a signature using the Notation CLI, you must configure the client so it can authenticate to Amazon ECR. If you have Docker installed on the same host where you install the Notation client, then Notation will reuse the same authentication method you use for the Docker client. The Docker `login` and `logout` commands will allow the Notation `sign` and `verify` commands to use those same credentials, and you don't have to separately authenticate Notation. For more information on configuring your Notation client for authentication, see Authenticate with OCI-compliant registries in the Notary Project documentation

If you are not using Docker or another tool that uses Docker credentials, then we recommend using the Amazon ECR Docker Credential Helper as your credential store. For more information on how to install and configure the Amazon ECR Credential Helper, see Amazon ECR Docker Credential Helper.

Signing an image

The following steps can be used to create the resources necessary to sign a container image and store the signature in an Amazon ECR private repository. Notation signs images using the digest.

To sign an image

1. Create an Amazon Signer signing profile using the `Notation-OCI-SHA384-ECDSA` signing platform. You can optionally specify a signature validity period using the `--signature-validity-period` parameter. This value may be specified using `DAYS`, `MONTHS`, or `YEARS`. If no validity period is specified, the default value of 135 months is used.

   ```
   aws signer put-signing-profile --profile-name ecr_signing_profile --platform-id
    Notation-OCI-SHA384-ECDSA
   ```

   > **Note**
   >
   > The signing profile name only supports alphanumeric characters and the underscore (_).

2. Authenticate the Notation client to your default registry. The following example uses the Amazon CLI to authenticate the Notation CLI to an Amazon ECR private registry.

```
aws ecr get-login-password --region region | notation login --username AWS --
password-stdin 111122223333.dkr.ecr.region.amazonaws.com
```

3. Use the Notation CLI to sign the image, specifying the image using the repository name and the SHA digest. This creates the signature and pushes it to the same Amazon ECR private repository that the image being signed is in.

   In the following example, we are signing an image in the `curl` repository with SHA digest `sha256:ca78e5f730f9a789ef8c63bb55275ac12dfb9e8099e6EXAMPLE`.

```
notation
 sign 111122223333.dkr.ecr.region.amazonaws.com/
curl@sha256:ca78e5f730f9a789ef8c63bb55275ac12dfb9e8099e6EXAMPLE --plugin
 "com.amazonaws.signer.notation.plugin" --id "arn:aws:signer:region:111122223333:/
signing-profiles/ecrSigningProfileName"
```

Next steps

After you sign your container image, you can verify the signature locally. For instructions about verifying an image, see Verify an image locally after signing in the *Amazon Signer Developer Guide*

Deleting a signature from an Amazon ECR private repository

You can delete a signature from an Amazon ECR private repository. When you create and push a signature using the Notation CLI, an OCI image index is created in your Amazon ECR repository as well. The Amazon ECR API doesn't support deleting artifacts or images referred to by an OCI image index, so the following are the available options to clean up these artifacts.

- (Recommended) You can use the ORAS CLI to delete the artifact and ORAS will handle updating or deleting the image index.
- You can use the Amazon ECR API or console to delete the OCI image index first and then the referenced artifact such as the signature.

When using the ORAS client to delete signatures and other reference type artifacts, ORAS manages the OCI image index. ORAS will first remove the reference to the artifact from the index, and then will delete the manifest. The `oras manifest delete` command can be used, referencing the index of the signature artifact.

To delete a signature using the ORAS CLI

1. Install and configure the ORAS client.

   For information about installing and configuring the ORAS client, see Installation in the ORAS documentation.

2. To delete a signature by using the ORAS CLI, run the following command:

```
oras manifest
 delete 111122223333.dkr.ecr.region.amazonaws.com/
repository_name@sha256:ca78e5f730f9a789ef8c63bb55275ac12dfb9e8099e6EXAMPLE
```

Viewing image details in Amazon ECR

After you push an image to your repository, you can view information about it. The details included are as follows:

- Image URI
- Image tags
- Artifact media type
- Image manifest type
- Scanning status
- The size of the image in MB
- When the image was pushed to the repository
- The replication status

To view image details (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region that contains the repository containing your image.
3. In the navigation pane, choose **Repositories**.
4. On the **Repositories** page, choose the repository to view.
5. On the **Repositories : *repository_name*** page, choose the image to view the details of.

Pulling an image to your local environment from an Amazon ECR private repository

If you want to run a Docker image that is available in Amazon ECR, you can pull it to your local environment with the **docker pull** command. You can do this from either your default registry or from a registry associated with another Amazon account.

To use an Amazon ECR image in an Amazon ECS task definition, see Using Amazon ECR images with Amazon ECS.

> **Important**
>
> Amazon ECR requires that users have permission to make calls to the `ecr:GetAuthorizationToken` API through an IAM policy before they can authenticate to a registry and push or pull any images from any Amazon ECR repository. Amazon ECR provides several Amazon managed policies to control user access at varying levels. For information about the Amazon managed policies for Amazon ECR, see Amazon managed policies for Amazon Elastic Container Registry.

To pull a Docker image from an Amazon ECR repository

1. Authenticate your Docker client to the Amazon ECR registry that you intend to pull your image from. Authentication tokens must be obtained for each registry used, and the tokens are valid for 12 hours. For more information, see Private registry authentication in Amazon ECR.
2. (Optional) Identify the image to pull.
   - You can list the repositories in a registry with the **aws ecr describe-repositories** command:

     ```
     aws ecr describe-repositories
     ```

     The example registry above has a repository called `amazonlinux`.
   - You can describe the images within a repository with the **aws ecr describe-images** command:

     ```
     aws ecr describe-images --repository-name amazonlinux
     ```

The example repository above has an image tagged as `latest` and `2016.09`, with the image digest `sha256:f1d4ae3f7261a72e98c6ebefe9985cf10a0ea5bd762585a43e0700ed99863807`.

3. Pull the image using the **docker pull** command. The image name format should be *`registry`*`/`*`repository`*`[:`*`tag`*`]` to pull by tag, or *`registry`*`/`*`repository`*`[@`*`digest`*`]` to pull by digest.

```
docker pull aws_account_id.dkr.ecr.us-west-2.amazonaws.com/amazonlinux:latest
```

> **Important**
>
> If you receive a *`repository-url`* `not found: does not exist or no pull access` error, you might need to authenticate your Docker client with Amazon ECR. For more information, see Private registry authentication in Amazon ECR.

Pulling the Amazon Linux container image

The Amazon Linux container image is built from the same software components that are included in the Amazon Linux AMI. The Amazon Linux container image is available for use in any environment as a base image for Docker workloads. If you use the Amazon Linux AMI for applications in Amazon EC2, you can containerize your applications with the Amazon Linux container image.

You can use the Amazon Linux container image in your local development environment and then push your application to Amazon using Amazon ECS. For more information, see Using Amazon ECR images with Amazon ECS.

The Amazon Linux container image is available on Docker Hub. For support for the Amazon Linux container image, go to the Amazon developer forums.

To pull the Amazon Linux container image from Docker Hub

1. Pull the Amazon Linux container image using the **docker pull** command.

```
docker pull amazonlinux
```

2. (Optional) Run the container locally.

```
docker run -it amazonlinux:latest /bin/bash
```

Deleting an image in Amazon ECR

If you're finished using an image, you can delete it from your repository. If you're finished with a repository, you can delete the entire repository and all of the images within it. For more information, see Deleting a private repository in Amazon ECR.

As an alternative to deleting images manually, you can create repository lifecycle policies which provide more control over the lifecycle management of images in your repositories. Lifecycle policies automate this process for you. For more information, see Automate the cleanup of images by using lifecycle policies in Amazon ECR.

To delete an image (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region that contains the image to delete.
3. In the navigation pane, choose **Repositories**.
4. On the **Repositories** page, choose the repository that contains the image to delete.
5. On the **Repositories:** ***repository_name*** page, select the box to the left of the image to delete and choose **Delete**.
6. In the **Delete image(s)** dialog box, verify that the selected images should be deleted and choose **Delete**.

To delete an image (Amazon CLI)

1. List the images in your repository. Tagged images will have both an image digest as well as a list of associated tags. Untagged images will only have an image digest.

   ```
   aws ecr list-images \
       --repository-name my-repo
   ```

2. (Optional) Delete any unwanted tags for the image by specifying the tag associated with the image you want to delete. When the last tag is deleted from an image, the image is also deleted.

```
aws ecr batch-delete-image \
     --repository-name my-repo \
     --image-ids imageTag=tag1 imageTag=tag2
```

3. Delete a tagged or untagged image by specifying the image digest. When you delete an image by referencing its digest, the image and all of its tags are deleted.

```
aws ecr batch-delete-image \
     --repository-name my-repo \
     --image-ids imageDigest=sha256:4f70ef7a4d29e8c0c302b13e25962d8f7a0bd304EXAMPLE
```

To delete multiple images, you can specify multiple image tags or image digests in the request.

```
aws ecr batch-delete-image \
     --repository-name my-repo \
     --image-ids imageDigest=sha256:4f70ef7a4d29e8c0c302b13e25962d8f7a0bd304EXAMPLE
 imageDigest=sha256:f5t0e245ssffc302b13e25962d8f7a0bd304EXAMPLE
```

Retagging an imagein Amazon ECR

With Docker Image Manifest V2 Schema 2 images, you can use the `--image-tag` option of the **put-image** command to retag an existing image. You can retag without pulling or pushing the image with Docker. For larger images, this process saves a considerable amount of network bandwidth and time required to retag an image.

To retag an image (Amazon CLI)

To retag an image with the Amazon CLI

1. Use the **batch-get-image** command to get the image manifest for the image to retag and write it to a file. In this example, the manifest for an image with the tag, *`latest`*, in the repository, *`amazonlinux`*, is written to an environment variable named *`MANIFEST`*.

```
MANIFEST=$(aws ecr batch-get-image --repository-name amazonlinux --image-ids
 imageTag=latest --output text --query 'images[].imageManifest')
```

2. Use the `--image-tag` option of the **put-image** command to put the image manifest to Amazon ECR with a new tag. In this example, the image is tagged as *`2017.03`*.

> **Note**
>
> If the `--image-tag` option isn't available in your version of the Amazon CLI, upgrade to the latest version. For more information, see Installing the Amazon Command Line Interface in the *Amazon Command Line Interface User Guide*.

```
aws ecr put-image --repository-name amazonlinux --image-tag 2017.03 --image-
manifest "$MANIFEST"
```

3. Verify that your new image tag is attached to your image. In the following output, the image has the tags `latest` and `2017.03`.

```
aws ecr describe-images --repository-name amazonlinux
```

The output is as follows:

```
{
    "imageDetails": [
        {
            "imageSizeInBytes": 98755613,
            "imageDigest":
 "sha256:8d00af8f076eb15a33019c2a3e7f1f655375681c4e5be157a26EXAMPLE",
            "imageTags": [
                "latest",
                "2017.03"
            ],
            "registryId": "aws_account_id",
            "repositoryName": "amazonlinux",
            "imagePushedAt": 1499287667.0
        }
    ]
}
```

To retag an image (Amazon Tools for Windows PowerShell)

To retag an image with the Amazon Tools for Windows PowerShell

1. Use the **Get-ECRImageBatch** cmdlet to obtain the description of the image to retag and write it to an environment variable. In this example, an image with the tag, *latest*, in the repository, *amazonlinux*, is written to the environment variable, *$Image*.

   > **Note**
   >
   > If you don't have the **Get-ECRImageBatch** cmdlet available on your system, see Setting up the Amazon Tools for Windows PowerShell in the *Amazon Tools for Windows PowerShell User Guide*.

   ```
   $Image = Get-ECRImageBatch -ImageId @{ imageTag="latest" } -
   RepositoryName amazonlinux
   ```

2. Write the manifest of the image to the *$Manifest* environment variable.

   ```
   $Manifest = $Image.Images[0].ImageManifest
   ```

3. Use the `-ImageTag` option of the **Write-ECRImage** cmdlet to put the image manifest to Amazon ECR with a new tag. In this example, the image is tagged as *2017.09*.

   ```
   Write-ECRImage -RepositoryName amazonlinux -ImageManifest $Manifest -
   ImageTag 2017.09
   ```

4. Verify that your new image tag is attached to your image. In the following output, the image has the tags `latest` and `2017.09`.

   ```
   Get-ECRImage -RepositoryName amazonlinux
   ```

   The output is as follows:

   ```
   ImageDigest                                                             ImageTag
   -----------                                                             --------
   sha256:359b948ea8866817e94765822787cd482279eed0c17bc674a7707f4256d5d497 latest
   sha256:359b948ea8866817e94765822787cd482279eed0c17bc674a7707f4256d5d497 2017.09
   ```

Preventing image tags from being overwritten in Amazon ECR

You can prevent image tags from being overwritten by turning on tag immutability in a repository. After tag immutability is turned on, the `ImageTagAlreadyExistsException` error is returned if you push an image with a tag that is already in the repository. Tag immutability affects all tags. You cannot make some tags immutable while others aren't.

You can use the Amazon Web Services Management Console and Amazon CLI tools to set image tag mutability for a new repository or for an existing repository. To create a repository using console steps, see Creating an Amazon ECR private repository to store images.

Setting image tag mutability (Amazon Web Services Management Console)

To set image tag mutability

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region that contains the repository to edit.
3. In the navigation pane, choose **Repositories**.
4. On the **Repositories** page, choose the **Private** tab and then select the repository to edit and choose **Edit**.
5. For **Tag immutability**, choose the tag mutability setting for the repository. Repositories configured with immutable tags prevent image tags from being overwritten. For more information, see Preventing image tags from being overwritten in Amazon ECR.
6. For **Image scan settings**, while you can specify the scan settings at the repository level for basic scanning, it is best practice to specify the scan configuration at the private registry level. Specify the scanning settings at the private registry allow you to enable either enhanced scanning or basic scanning as well as define filters to specify which repositories are scanned. For more information, see Scan images for software vulnerabilities in Amazon ECR.
7. For **Encryption settings**, this is a view only field as the encryption settings for a repository can't be changed once the repository is created.
8. Choose **Save** to update the repository settings.

Setting image tag mutability (Amazon CLI)

To create a repository with immutable tags configured

Use one of the following commands to create a new image repository with immutable tags configured.

- create-repository (Amazon CLI)

```
aws ecr create-repository --repository-name name --image-tag-mutability IMMUTABLE --
region us-east-2
```

- New-ECRRepository (Amazon Tools for Windows PowerShell)

```
New-ECRRepository -RepositoryName name -ImageTagMutability IMMUTABLE -Region us-
east-2 -Force
```

To update the image tag mutability settings for a repository

Use one of the following commands to update the image tag mutability settings for an existing repository.

- put-image-tag-mutability (Amazon CLI)

```
aws ecr put-image-tag-mutability --repository-name name --image-tag-
mutability IMMUTABLE --region us-east-2
```

- Write-ECRImageTagMutability (Amazon Tools for Windows PowerShell)

```
Write-ECRImageTagMutability -RepositoryName name -ImageTagMutability IMMUTABLE -
Region us-east-2 -Force
```

Container image manifest format support in Amazon ECR

Amazon ECR supports the following container image manifest formats:

- Docker Image Manifest V2 Schema 1 (used with Docker version 1.9 and older)
- Docker Image Manifest V2 Schema 2 (used with Docker version 1.10 and newer)

- Open Container Initiative (OCI) Specifications (v1.0 and up)

Support for Docker Image Manifest V2 Schema 2 provides the following functionality:

- The ability to use multiple tags for a singular image.
- Support for storing Windows container images.

Amazon ECR image manifest conversion

When you push and pull images to and from Amazon ECR, your container engine client (for example, Docker) communicates with the registry to agree on a manifest format that is understood by the client and the registry to use for the image.

When you push an image to Amazon ECR with Docker version 1.9 or earlier, the image manifest format is stored as Docker Image Manifest V2 Schema 1. When you push an image to Amazon ECR with Docker version 1.10 or later, the image manifest format is stored as Docker Image Manifest V2 Schema 2.

When you pull an image from Amazon ECR *by tag*, Amazon ECR returns the image manifest format that is stored in the repository. The format is returned only if that format is understood by the client. If the stored image manifest format isn't understood by the client, Amazon ECR converts the image manifest into a format that is understood. For example, if a Docker 1.9 client requests an image manifest that is stored as Docker Image Manifest V2 Schema 2, Amazon ECR returns the manifest in the Docker Image Manifest V2 Schema 1 format. The following table describes the available conversions supported by Amazon ECR when an image is pulled *by tag*:

Schema requested by client	Pushed to ECR as V2, schema 1	Pushed to ECR as V2, schema 2	Pushed to ECR as OCI
V2, schema 1	No translation required	Translated to V2, schema 1	Translated to V2, schema 1
V2, schema 2	No translation available, client falls back to V2, schema 1	No translation required	Translated to V2, schema 2

Schema requested by client	Pushed to ECR as V2, schema 1	Pushed to ECR as V2, schema 2	Pushed to ECR as OCI
OCI	No translation available	Translated to OCI	No translation required

Important

If you pull an image *by digest*, there is no translation available. Your client must understand the image manifest format that is stored in Amazon ECR. If you request a Docker Image Manifest V2 Schema 2 image by digest on a Docker 1.9 or older client, the image pull fails. For more information, see Registry compatibility in the Docker documentation.
In this example, if you request the same image *by tag*, Amazon ECR translates the image manifest into a format that the client can understand. The image pull succeeds.

Using Amazon ECR images with Amazon ECS

You can use your Amazon ECR private repositories to host container images and artifacts that your Amazon ECS tasks may pull from. For this to work, the Amazon ECS, or Fargate, container agent must have permissions to make the `ecr:BatchGetImage`, `ecr:GetDownloadUrlForLayer`, and `ecr:GetAuthorizationToken` APIs.

Required IAM permissions

The following table shows the IAM role to use, for each launch type, that provides the required permissions for your tasks to pull from an Amazon ECR private repository. Amazon ECS provides managed IAM policies that include the required permissions.

Launch type	IAM role	Amazon managed IAM policy
Amazon ECS on Amazon EC2 instances	Use the container instance IAM role, which is associate d with the Amazon EC2 instance registered to your Amazon ECS cluster. For more	`AmazonEC2Container ServiceforEC2Role` For more information, see AmazonEC2Container ServiceforEC2Role in the

Launch type	IAM role	Amazon managed IAM policy
	information, see Container instance IAM role in the *Amazon Elastic Container Service Developer Guide*.	*Amazon Elastic Container Service Developer Guide*
Amazon ECS on Fargate	Use the task execution IAM role that you reference in your Amazon ECS task definition. For more information, see Task execution IAM role in the *Amazon Elastic Container Service Developer Guide*.	`AmazonECSTaskExecutionRolePolicy` For more information, see AmazonECSTaskExecutionRolePolicy in the *Amazon Elastic Container Service Developer Guide*.
Amazon ECS on external instances	Use the container instance IAM role, which is associated with the on-premises server or virtual machine (VM) registered to your Amazon ECS cluster. For more information, see Container instance Amazon ECS role in the *Amazon Elastic Container Service Developer Guide*.	`AmazonEC2ContainerServiceforEC2Role` For more information, see AmazonEC2ContainerServiceforEC2Role in the *Amazon Elastic Container Service Developer Guide*.

Important

The Amazon managed IAM policies contain additional permissions that you may not require for your use. In this case, these are the minimum required permissions to pull from an Amazon ECR private repository.

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

```
            "Action": [
                "ecr:BatchGetImage",
                "ecr:GetDownloadUrlForLayer",
                "ecr:GetAuthorizationToken"
            ],
            "Resource": "*"
        }
    ]
}
```

Specifying an Amazon ECR image in an Amazon ECS task definition

When creating an Amazon ECS task definition, you can specify a container image hosted in an Amazon ECR private repository. In the task definition, ensure that you use the full `registry/repository:tag` naming for your Amazon ECR images. For example, *aws_account_id*.dkr.ecr.*region*.amazonaws.com/*my-repository*:*latest*.

The following task definition snippet shows the syntax you would use to specify a container image hosted in Amazon ECR in your Amazon ECS task definition.

```
{
    "family": "task-definition-name",
    ...
    "containerDefinitions": [
        {
            "name": "container-name",
            "image": "aws_account_id.dkr.ecr.region.amazonaws.com/my-
repository:latest",
            ...
        }
    ],
    ...
}
```

Using Amazon ECR Images with Amazon EKS

You can use your Amazon ECR images with Amazon EKS.

When referencing an image from Amazon ECR, you must use the full `registry/repository:tag` naming for the image. For example, *`aws_account_id`*`.dkr.ecr.`*`region`*`.amazonaws.com/`*`my-repository`*`:`*`latest`*.

Required IAM permissions

If you have Amazon EKS workloads hosted on managed nodes, self-managed nodes, or Amazon Fargate, review the following:

- Amazon EKS workloads hosted on managed or self-managed nodes: The Amazon EKS worker node IAM role (`NodeInstanceRole`) is required. The Amazon EKS worker node IAM role must contain the following IAM policy permissions for Amazon ECR.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "ecr:BatchCheckLayerAvailability",
                "ecr:BatchGetImage",
                "ecr:GetDownloadUrlForLayer",
                "ecr:GetAuthorizationToken"
            ],
            "Resource": "*"
        }
    ]
}
```

> **Note**
>
> If you used `eksctl` or the Amazon CloudFormation templates in Getting Started with Amazon EKS to create your cluster and worker node groups, these IAM permissions are applied to your worker node IAM role by default.

- Amazon EKS workloads hosted on Amazon Fargate: Use the Fargate pod execution role, which provides your pods permission to pull images from private Amazon ECR repositories. For more information, see Create a Fargate pod execution role.

Installing a Helm chart on an Amazon EKS cluster

Helm charts hosted in Amazon ECR can be installed on your Amazon EKS clusters.

Prerequisites

- Install the latest version of the Helm client. These steps were written using Helm version `3.9.0`. For more information, see Installing Helm.
- You have at least version `1.23.9` or `2.6.3` of the Amazon CLI installed on your computer. For more information, see Installing or updating the latest version of the Amazon CLI.
- You have pushed a Helm chart to your Amazon ECR repository. For more information, see Pushing a Helm chart to an Amazon ECR private repository.
- You have configured `kubectl` to work with Amazon EKS. For more information, see Create a `kubeconfig` for Amazon EKS in the Amazon EKS User Guide. If the following commands succeeds for your cluster, you're properly configured.

```
kubectl get svc
```

To install a Helm chart on an Amazon EKS cluster

1. Authenticate your Helm client to the Amazon ECR registry that your Helm chart is hosted. Authentication tokens must be obtained for each registry used, and the tokens are valid for 12 hours. For more information, see Private registry authentication in Amazon ECR.

```
aws ecr get-login-password \
     --region us-west-2 | helm registry login \
     --username AWS \
     --password-stdin aws_account_id.dkr.ecr.region.amazonaws.com
```

2. Install the chart. Replace *`helm-test-chart`* with your repository and *`0.1.0`* with your Helm chart's tag.

```
helm install ecr-chart-demo oci://aws_account_id.dkr.ecr.region.amazonaws.com/helm-
test-chart --version 0.1.0
```

The output should look similar to this:

```
NAME: ecr-chart-demo
```

```
LAST DEPLOYED: Tue May 31 17:38:56 2022
NAMESPACE: default
STATUS: deployed
REVISION: 1
TEST SUITE: None
```

3. Verify the chart installation.

```
helm list -n default
```

Example output:

```
NAME            NAMESPACE       REVISION        UPDATED
     STATUS          CHART                   APP VERSION
ecr-chart-demo  default         1               2022-06-01 15:56:40.128669157 +0000
 UTC deployed        helm-test-chart-0.1.0   1.16.0
```

4. (Optional) See the installed Helm chart `ConfigMap`.

```
kubectl describe configmap helm-test-chart-configmap
```

5. When you are finished, you can remove the chart release from your cluster.

```
helm uninstall ecr-chart-demo
```

Scan images for software vulnerabilities in Amazon ECR

The improved basic scanning feature is in preview release for Amazon ECR and is subject to change. During this public preview, you can only use the Amazon Web Services Management Console to opt-in for the Improved basic scanning version.

Amazon ECR image scanning helps to identify software vulnerabilities in your container images. The following scanning types are offered.

> **Important**
>
> Switching between the Enhanced scanning, Basic scanning, and the Improved basic scanning versions will cause previously established scans to no longer be available. You will have to set up your scans again. However, if you switch back to your previous scanning version the established scans will be available.

- **Enhanced scanning**—Amazon ECR integrates with Amazon Inspector to provide automated, continuous scanning of your repositories. Your container images are scanned for both operating systems and programing language package vulnerabilities. As new vulnerabilities appear, the scan results are updated and Amazon Inspector emits an event to EventBridge to notify you. Enhanced scanning provides the following:
  - OS and programming languages package vulnerabilities.
  - Two scanning frequencies: Scan on push and continuous scan.
- **Basic scanning**—Amazon ECR provides two versions of basic scanning which use the Common Vulnerabilities and Exposures (CVEs) database: the current GA version that uses the open-source Clair project and a newly improved version of basic scanning (in preview) that uses our Amazon native technology. With basic scanning, you configure your repositories to scan on push or you can perform manual scans and Amazon ECR provides a list of scan findings. Basic scanning provides the following:
  - OS scans.
  - Two scanning frequencies: Manual and scan on push.

> **Important**
>
> The new version of basic scanning does not support `imageScanFindingsSummary` and `imageScanStatus` in the `DescribeImages` API. To view these, use the `DescribeImageScanFindings` API.

Filters to choose which repositories are scanned in Amazon ECR

When you configure image scanning for your private registry, you can use filters to choose which repositories are scanned.

When **basic** scanning is used, you may specify scan on push filters to specify which repositories are set to do an image scan when new images are pushed. Any repositories not matching a basic scanning scan on push filter will be set to the **manual** scan frequency which means to perform a scan, you must manually trigger the scan.

When **enhanced** scanning is used, you may specify separate filters for scan on push and continuous scanning. Any repositories not matching an enhanced scanning filter will have scanning disabled. If you are using enhanced scanning and specify separate filters for scan on push and continuous scanning where multiple filters match the same repository, then Amazon ECR enforces the continuous scanning filter over the scan on push filter for that repository.

Filter wildcards

When a filter is specified, a filter with no wildcard will match all repository names that contain the filter. A filter with a wildcard (*) matches on any repository name where the wildcard replaces zero or more characters in the repository name.

The following table provides examples where repository names are expressed on the horizontal axis and example filters are specified on the vertical axis.

	prod	**repo-prod**	**prod-repo**	**repo-prod-repo**	**prodrepo**
prod	Yes	Yes	Yes	Yes	Yes
*prod	Yes	Yes	No	No	No

	prod	repo-prod	prod-repo	repo-prod-repo	prodrepo
prod*	Yes	No	Yes	No	Yes
prod	Yes	Yes	Yes	Yes	Yes
prod*repo	No	No	Yes	No	Yes

Scan images for OS and programming language package vulnerabilities in Amazon ECR

Amazon ECR enhanced scanning is an integration with Amazon Inspector which provides vulnerability scanning for your container images. Your container images are scanned for both operating systems and programming language package vulnerabilities. You can view the scan findings with both Amazon ECR and with Amazon Inspector directly. For more information about Amazon Inspector, see Scanning container images with Amazon Inspector in the *Amazon Inspector User Guide*.

With enhanced scanning, you can choose which repositories are configured for automatic, continuous scanning and which are configured for scan on push. This is done by setting scan filters.

Considerations for enhanced scanning

Consider the following before enabling Amazon ECR enhanced scanning.

- There is no additional cost from Amazon ECR to use this feature, however there is a cost from Amazon Inspector to scan your images. For more information, see Amazon Inspector pricing.
- Enhanced scanning isn't supported in the following Regions:
  - Middle East (UAE) (`me-central-1`)
  - Asia Pacific (Hyderabad) (`ap-south-2`)
  - Israel (Tel Aviv) (`il-central-1`)
  - Asia Pacific (Melbourne) (`ap-southeast-4`)
  - Europe (Spain) (`eu-south-2`)
- Amazon Inspector supports scanning for specific operating systems. For a full list, see Supported operating systems - Amazon ECR scanning in the *Amazon Inspector User Guide*.

- Amazon Inspector uses a service-linked IAM role, which provides the permissions needed to provide enhanced scanning for your repositories. The service-linked IAM role is created automatically by Amazon Inspector when enhanced scanning is turned on for your private registry. For more information, see Using service-linked roles for Amazon Inspector in the *Amazon Inspector User Guide*.
- When you initially turn on enhanced scanning for your private registry, Amazon Inspector only recognizes images pushed to Amazon ECR in the last 30 days, based on the image push timestamp, or pulled in the last 90 days. Older images will have the `SCAN_ELIGIBILITY_EXPIRED` scan status. If you'd like these images to be scanned by Amazon Inspector you should push them again to your repository.
- All images pushed to Amazon ECR after enhanced scanning is turned on are continually scanned for the configured duration. By default, the duration is **Lifetime**. This setting can be configured using the Amazon Inspector console. For more information, see Changing the enhanced scanning duration for images in Amazon Inspector.
- When enhanced scanning is turned on for your Amazon ECR private registry, repositories matching the scan filters are scanned using enhanced scanning only. Any repositories that don't match a filter will have an `Off` scan frequency and won't be scanned. Manual scans using enhanced scanning aren't supported. For more information, see Filters to choose which repositories are scanned in Amazon ECR.
- If you specify separate filters for scan on push and continuous scanning where multiple filters match the same repository, then Amazon ECR enforces the continuous scanning filter over the scan on push filter for that repository.
- When enhanced scanning is turned on, Amazon ECR sends an event to EventBridge when the scan frequency for a repository is changed. Amazon Inspector emits events to EventBridge when an initial scan is completed and when an image scan finding is created, updated, or closed.

IAM permissions required for enhanced scanning in Amazon ECR

Amazon ECR enhanced scanning requires an Amazon Inspector service-linked IAM role and that the IAM principal enabling and using enhanced scanning has permissions to call the Amazon Inspector APIs needed for scanning. The Amazon Inspector service-linked IAM role is created automatically by Amazon Inspector when enhanced scanning is turned on for your private registry. For more information, see Using service-linked roles for Amazon Inspector in the *Amazon Inspector User Guide*.

The following IAM policy grants the required permissions for enabling and using enhanced scanning. It includes the permission needed for Amazon Inspector to create the service-linked IAM role as well as the Amazon Inspector API permissions needed to turned on and off enhanced scanning and retrieve the scan findings.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "inspector2:Enable",
                "inspector2:Disable",
                "inspector2:ListFindings",
                "inspector2:ListAccountPermissions",
                "inspector2:ListCoverage"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": "iam:CreateServiceLinkedRole",
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "iam:AWSServiceName": [
                        "inspector2.amazonaws.com"
                    ]
                }
            }
        }
    ]
}
```

Configuring enhanced scanning for images in Amazon ECR

Configure enhanced scanning per Region for your private registry.

Verify that you have the proper IAM permissions to configure enhanced scanning. For information, see IAM permissions required for enhanced scanning in Amazon ECR.

Amazon Web Services Management Console

To turn on enhanced scanning for your private registry

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region to set the scanning configuration for.
3. In the navigation pane, choose **Private registry**, **Settings**, **Scanning**.
4. On the **Scanning configuration** page, for **Scan type** choose **Enhanced scanning**.

   By default, when **Enhanced scanning** is selected, all of your repositories are continuously scanned.
5. To choose specific repositories to continuously scan, clear the **Continuously scan all repositories** box, and then define your filters:

   > **Important**
   >
   > Filters with no wildcard will match all repository names that contain the filter. Filters with wildcards (*) match on a repository name where the wildcard replaces zero or more characters in the repository name. To see examples of how filters behave, see the section called "Filter wildcards".

   a. Enter a filter based on repository names, and then choose **Add filter**.
   b. Decide which repositories to scan when an image is pushed:
      - To scan all repositories on push, select **Scan on push all repositories**.
      - To choose specific repositories to scan on push, enter a filter based on repository names, and then choose **Add filter**.
6. Choose **Save**.
7. Repeat these steps in each Region in which you want to turn on enhanced scanning.

Amazon CLI

Use the following Amazon CLI command to turn on enhanced scanning for your private registry using the Amazon CLI. You can specify scan filters using the `rules` object.

- put-registry-scanning-configuration (Amazon CLI)

The following example turns on enhanced scanning for your private registry. By default, when no `rules` are specified, Amazon ECR sets the scanning configuration to continuous scanning for all repositories.

```
aws ecr put-registry-scanning-configuration \
     --scan-type ENHANCED \
     --region us-east-2
```

The following example turns on enhanced scanning for your private registry and specifies a scan filter. The scan filter in the example turns on continuous scanning for all repositories with `prod` in its name.

```
aws ecr put-registry-scanning-configuration \
     --scan-type ENHANCED \
     --rules '[{"repositoryFilters" : [{"filter":"prod","filterType" :
 "WILDCARD"}],"scanFrequency" : "CONTINUOUS_SCAN"}]' \
     --region us-east-2
```

The following example turns on enhanced scanning for your private registry and specifies multiple scan filters. The scan filters in the example turns on continuous scanning for all repositories with `prod` in its name and turns on scan on push only for all other repositories.

```
aws ecr put-registry-scanning-configuration \
     --scan-type ENHANCED \
     --rules '[{"repositoryFilters" : [{"filter":"prod","filterType" :
 "WILDCARD"}],"scanFrequency" : "CONTINUOUS_SCAN"},{"repositoryFilters" :
 [{"filter":"*","filterType" : "WILDCARD"}],"scanFrequency" : "SCAN_ON_PUSH"}]' \
     --region us-west-2
```

Changing the enhanced scanning duration for images in Amazon Inspector

You can change the number of days that Amazon Inspector continuously scans the images in your Amazon ECR private repositories. By default, when enhanced scanning is turned on for your Amazon ECR private registry, the Amazon Inspector service continually monitors your repositories until either the image is deleted or enhanced scanning is disabled. The duration that Amazon Inspector scans your images can be changed using the Amazon Inspector settings. The available

scan durations are **Lifetime (default)**, **180 days**, and **30 days**. When the scan duration for a repository elapses, the scan status of `SCAN_ELIGIBILITY_EXPIRED` is displayed when listing your scan vulnerabilities. For more information, see Changing the Amazon ECR automated re-scan duration in the *Amazon Inspector User Guide*.

To change the enhanced scanning duration setting

1. Open the Amazon Inspector console at https://console.aws.amazon.com/inspector/v2/home.
2. In the left navigation, expand **Settings** and then choose **General**.
3. On the **Settings** page, under **ECR re-scan duration** choose a setting, then choose **Save**.

EventBridge events sent for enhanced scanning in Amazon ECR

When enhanced scanning is turned on, Amazon ECR sends an event to EventBridge when the scan frequency for a repository is changed. Amazon Inspector sends events to EventBridge when an initial scan is completed and when an image scan finding is created, updated, or closed.

Event for a repository scan frequency change

When enhanced scanning is turned on for your registry, the following event is sent by Amazon ECR when there is a change with a resource that has enhanced scanning turned on. This includes new repositories being created, the scan frequency for a repository being changed, or when images are created or deleted in repositories with enhanced scanning turned on. For more information, see Scan images for software vulnerabilities in Amazon ECR.

```
{
 "version": "0",
 "id": "0c18352a-a4d4-6853-ef53-0abEXAMPLE",
 "detail-type": "ECR Scan Resource Change",
 "source": "aws.ecr",
 "account": "123456789012",
 "time": "2021-10-14T20:53:46Z",
 "region": "us-east-1",
 "resources": [],
 "detail": {
  "action-type": "SCAN_FREQUENCY_CHANGE",
  "repositories": [{
    "repository-name": "repository-1",
    "repository-arn": "arn:aws:ecr:us-east-1:123456789012:repository/repository-1",
    "scan-frequency": "SCAN_ON_PUSH",
```

```
    "previous-scan-frequency": "MANUAL"
   },
   {
    "repository-name": "repository-2",
    "repository-arn": "arn:aws:ecr:us-east-1:123456789012:repository/repository-2",
    "scan-frequency": "CONTINUOUS_SCAN",
    "previous-scan-frequency": "SCAN_ON_PUSH"
   },
   {
    "repository-name": "repository-3",
    "repository-arn": "arn:aws:ecr:us-east-1:123456789012:repository/repository-3",
    "scan-frequency": "CONTINUOUS_SCAN",
    "previous-scan-frequency": "SCAN_ON_PUSH"
   }
  ],
  "resource-type": "REPOSITORY",
  "scan-type": "ENHANCED"
 }
}
```

Event for an initial image scan (enhanced scanning)

When enhanced scanning is turned on for your registry, the following event is sent by Amazon Inspector when the initial image scan is completed. The `finding-severity-counts` parameter will only return a value for a severity level if one exists. For example, if the image contains no findings at `CRITICAL` level, then no critical count is returned. For more information, see Scan images for OS and programming language package vulnerabilities in Amazon ECR.

Event pattern:

```
{
  "source": ["aws.inspector2"],
  "detail-type": ["Inspector2 Scan"]
}
```

Example output:

```
{
    "version": "0",
    "id": "739c0d3c-4f02-85c7-5a88-94a9EXAMPLE",
    "detail-type": "Inspector2 Scan",
    "source": "aws.inspector2",
```

```
    "account": "123456789012",
    "time": "2021-12-03T18:03:16Z",
    "region": "us-east-2",
    "resources": [
        "arn:aws:ecr:us-east-2:123456789012:repository/amazon/amazon-ecs-sample"
    ],
    "detail": {
        "scan-status": "INITIAL_SCAN_COMPLETE",
        "repository-name": "arn:aws:ecr:us-east-2:123456789012:repository/amazon/
amazon-ecs-sample",
        "finding-severity-counts": {
            "CRITICAL": 7,
            "HIGH": 61,
            "MEDIUM": 62,
            "TOTAL": 158
        },
        "image-digest":
 "sha256:36c7b282abd0186e01419f2e58743e1bf635808231049bbc9d77e5EXAMPLE",
        "image-tags": [
            "latest"
        ]
    }
}
```

Event for an image scan finding update (enhanced scanning)

When enhanced scanning is turned on for your registry, the following event is sent by Amazon Inspector when the image scan finding is created, updated, or closed. For more information, see Scan images for OS and programming language package vulnerabilities in Amazon ECR.

Event pattern:

```
{
  "source": ["aws.inspector2"],
  "detail-type": ["Inspector2 Finding"]
}
```

Example output:

```
{
    "version": "0",
    "id": "42dbea55-45ad-b2b4-87a8-afaEXAMPLE",
    "detail-type": "Inspector2 Finding",
```

```
    "source": "aws.inspector2",
    "account": "123456789012",
    "time": "2021-12-03T18:02:30Z",
    "region": "us-east-2",
    "resources": [
        "arn:aws:ecr:us-east-2:123456789012:repository/amazon/amazon-ecs-sample/
sha256:36c7b282abd0186e01419f2e58743e1bf635808231049bbc9d77eEXAMPLE"
    ],
    "detail": {
        "awsAccountId": "123456789012",
        "description": "In libssh2 v1.9.0 and earlier versions, the SSH_MSG_DISCONNECT
 logic in packet.c has an integer overflow in a bounds check, enabling an attacker to
 specify an arbitrary (out-of-bounds) offset for a subsequent memory read. A crafted
 SSH server may be able to disclose sensitive information or cause a denial of service
 condition on the client system when a user connects to the server.",
        "findingArn": "arn:aws:inspector2:us-east-2:123456789012:finding/
be674aaddd0f75ac632055EXAMPLE",
        "firstObservedAt": "Dec 3, 2021, 6:02:30 PM",
        "inspectorScore": 6.5,
        "inspectorScoreDetails": {
            "adjustedCvss": {
                "adjustments": [],
                "cvssSource": "REDHAT_CVE",
                "score": 6.5,
                "scoreSource": "REDHAT_CVE",
                "scoringVector": "CVSS:3.0/AV:N/AC:L/PR:N/UI:R/S:U/C:H/I:N/A:N",
                "version": "3.0"
            }
        },
        "lastObservedAt": "Dec 3, 2021, 6:02:30 PM",
        "packageVulnerabilityDetails": {
            "cvss": [
                {
                    "baseScore": 6.5,
                    "scoringVector": "CVSS:3.0/AV:N/AC:L/PR:N/UI:R/S:U/C:H/I:N/A:N",
                    "source": "REDHAT_CVE",
                    "version": "3.0"
                },
                {
                    "baseScore": 5.8,
                    "scoringVector": "AV:N/AC:M/Au:N/C:P/I:N/A:P",
                    "source": "NVD",
                    "version": "2.0"
                },
```

```
                {
                    "baseScore": 8.1,
                    "scoringVector": "CVSS:3.1/AV:N/AC:L/PR:N/UI:R/S:U/C:H/I:N/A:H",
                    "source": "NVD",
                    "version": "3.1"
                }
            ],
            "referenceUrls": [
                "https://access.redhat.com/errata/RHSA-2020:3915"
            ],
            "source": "REDHAT_CVE",
            "sourceUrl": "https://access.redhat.com/security/cve/CVE-2019-17498",
            "vendorCreatedAt": "Oct 16, 2019, 12:00:00 AM",
            "vendorSeverity": "Moderate",
            "vulnerabilityId": "CVE-2019-17498",
            "vulnerablePackages": [
                {
                    "arch": "X86_64",
                    "epoch": 0,
                    "name": "libssh2",
                    "packageManager": "OS",
                    "release": "12.amzn2.2",
                    "sourceLayerHash":
"sha256:72d97abdfae3b3c933ff41e39779cc72853d7bd9dc1e4800c5294dEXAMPLE",
                    "version": "1.4.3"
                }
            ]
        },
        "remediation": {
            "recommendation": {
                "text": "Update all packages in the vulnerable packages section to
their latest versions."
            }
        },
        "resources": [
            {
                "details": {
                    "awsEcrContainerImage": {
                        "architecture": "amd64",
                        "imageHash":
"sha256:36c7b282abd0186e01419f2e58743e1bf635808231049bbc9d77e5EXAMPLE",
                        "imageTags": [
                            "latest"
                        ],
```

```
                        "platform": "AMAZON_LINUX_2",
                        "pushedAt": "Dec 3, 2021, 6:02:13 PM",
                        "registry": "123456789012",
                        "repositoryName": "amazon/amazon-ecs-sample"
                    }
                },
                "id": "arn:aws:ecr:us-east-2:123456789012:repository/amazon/amazon-ecs-
sample/sha256:36c7b282abd0186e01419f2e58743e1bf635808231049bbc9d77EXAMPLE",
                "partition": "N/A",
                "region": "N/A",
                "type": "AWS_ECR_CONTAINER_IMAGE"
            }
        ],
        "severity": "MEDIUM",
        "status": "ACTIVE",
        "title": "CVE-2019-17498 - libssh2",
        "type": "PACKAGE_VULNERABILITY",
        "updatedAt": "Dec 3, 2021, 6:02:30 PM"
    }
}
```

Retrieving the findings for enhanced scans in Amazon ECR

You can retrieve the scan findings for the last completed enhanced image scan, and then open the findings in Amazon Inspector to see more detail. The software vulnerabilities that were discovered are listed by severity based on the Common Vulnerabilities and Exposures (CVEs) database.

For troubleshooting details for some common issues when scanning images, see Troubleshooting image scanning in Amazon ECR.

Amazon Web Services Management Console

Use the following steps to retrieve image scan findings using the Amazon Web Services Management Console.

To retrieve image scan findings

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region where your repository exists.
3. In the navigation pane, choose **Repositories**.

4. On the **Repositories** page, choose the repository that contains the image to retrieve the scan findings for.
5. On the **Images** page, under the **Vulnerabilities** column, select **See findings** for the image to retrieve the scan findings for.
6. To view more details in the Amazon Inspector console, choose the vulnerability name in the **Name** column.

Amazon CLI

Use the following Amazon CLI command to retrieve image scan findings using the Amazon CLI. You can specify an image using the `imageTag` or `imageDigest`, both of which can be obtained using the list-images CLI command.

- describe-image-scan-findings (Amazon CLI)

  The following example uses an image tag.

```
aws ecr describe-image-scan-findings \
     --repository-name name \
     --image-id imageTag=tag_name \
     --region us-east-2
```

  The following example uses an image digest.

```
aws ecr describe-image-scan-findings \
     --repository-name name \
     --image-id imageDigest=sha256_hash \
     --region us-east-2
```

Scan images for OS vulnerabilities in Amazon ECR

The improved basic scanning feature is in preview release for Amazon ECR and is subject to change. During this public preview, you can only use the Amazon Web Services Management Console to opt-in for the Improved basic scanning version.

Amazon ECR provides two versions of basic scanning that use the Common Vulnerabilities and Exposures (CVEs) database:

- The current GA version that uses the open-source Clair project. For more information about Clair, see Clair on GitHub.
- The newly improved version of basic scanning (in preview) that uses Amazon native technology.

Amazon ECR uses the severity for a CVE from the upstream distribution source if available. Otherwise, the Common Vulnerability Scoring System (CVSS) score is used. The CVSS score can be used to obtain the NVD vulnerability severity rating. For more information, see NVD Vulnerability Severity Ratings.

Both versions of Amazon ECR basic scanning support filters to specify which repositories to scan on push. Any repositories that don't match a scan on push filter are set to the **manual** scan frequency which means you must manually start the scan. An image can be scanned once per 24 hours. The 24 hours includes the initial scan on push, if configured, and any manual scans.

The last completed image scan findings can be retrieved for each image. When an image scan is completed, Amazon ECR sends an event to Amazon EventBridge. For more information, see Amazon ECR events and EventBridge.

Region support for improved basic scanning

The improved version of basic scanning is supported in the following regions:

- Asia Pacific (Hong Kong) (`ap-east-1`)
- Europe (Stockholm) (`eu-north-1`)
- Middle East (Bahrain) (`me-south-1`)
- Asia Pacific (Mumbai) (`ap-south-1`)
- Europe (Paris) (`eu-west-3`)
- Amazon GovCloud (US-East) (`us-gov-east-1`)
- Africa (Cape Town) (`af-south-1`)
- Asia Pacific (Jakarta) (`ap-southeast-3`)
- Europe (Frankfurt) (`eu-central-1`)
- Europe (Ireland) (`eu-west-1`)
- South America (São Paulo) (`sa-east-1`)

- US East (Ohio) (`us-east-2`)
- Amazon GovCloud (US-West) (`us-gov-west-1`)
- Asia Pacific (Tokyo) (`ap-northeast-1`)
- Asia Pacific (Seoul) (`ap-northeast-2`)
- Asia Pacific (Osaka) (`ap-northeast-3`)
- Europe (Milan) (`eu-south-1`)
- Europe (London) (`eu-west-2`)
- US East (N. Virginia) (`us-east-1`)
- Asia Pacific (Singapore) (`ap-southeast-1`)
- Asia Pacific (Sydney) (`ap-southeast-2`)
- Canada (Central) (`ca-central-1`)
- US West (N. California) (`us-west-1`)
- US West (Oregon) (`us-west-2`)
- Europe (Zurich) (`eu-central-2`)

Operating system support for basic scanning and improved basic scanning

As a security best practice and for continued coverage, we recommend that you continue to use supported versions of an operating system. In accordance with vendor policy, discontinued operating systems are no longer updated with patches and, in many cases, new security advisories are no longer released for them. In addition, some vendors remove existing security advisories and detections from their feeds when an affected operating system reaches the end of standard support. After a distribution loses support from its vendor, Amazon ECR may no longer support scanning it for vulnerabilities. Any findings that Amazon ECR does generate for a discontinued operating system should be used for informational purposes only. Listed below are the current supported operating systems and versions.

Operating System	Version
Alpine Linux (Alpine)	3.19
Alpine Linux (Alpine)	3.18

Operating System	Version
Alpine Linux (Alpine)	3.17
Alpine Linux (Alpine)	3.16
Amazon Linux 2 (AL2)	AL2
Amazon Linux 2023(AL2023)	AL2023
CentOS Linux (CentOS)	7
Debian Server (Bookworm)	12
Debian Server (Bullseye)	11
Debian Server (Buster)	10
Oracle Linux (Oracle)	9
Oracle Linux (Oracle)	8
Oracle Linux (Oracle)	7
Ubuntu (Lunar)	23.04
Ubuntu (Jammy)	22.04 (LTS)
Ubuntu (Focal)	20.024 (LTS)
Ubuntu (Bionic)	18.04 (ESM)
Ubuntu (Xenial)	16.04 (ESM)
Ubuntu (Trusty)	14.04 (ESM)
Red Hat Enterprise Linux (RHEL)	7
Red Hat Enterprise Linux (RHEL)	8
Red Hat Enterprise Linux (RHEL)	9

Configuring improved basic scanning for images in Amazon ECR

An improved version of Amazon ECR basic scanning is now available in preview. Improved basic scanning uses Amazon native technology.

Configure improved basic scanning per Region for your private repository. For a list of regions that support improved basic scanning, see Region support for improved basic scanning.

To turn on improved basic scanning for your private registry

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region to set the scanning configuration for.
3. In the navigation pane, choose **Private registry**, **Scanning**.
4. On the **Scanning configuration** page, For **Scan type** choose **Improved basic scanning (In preview) - new**.
5. By default all of your repositories are set for **Manual** scanning. You can optionally configure scan on push by specifying **Scan on push filters**. You can set scan on push for all repositories or individual repositories. For more information, see Filters to choose which repositories are scanned in Amazon ECR.

Configuring basic scanning for images in Amazon ECR

By default, Amazon ECR turns on basic scanning for all private registries. As a result, unless you've changed the scanning settings on your private registry there is no need to turn on basic scanning. Basic scanning uses the open-source Clair project.

You can use the following steps to define one or more scan on push filters.

To turn on basic scanning for your private registry

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region to set the scanning configuration for.
3. In the navigation pane, choose **Private registry**, **Scanning**.
4. On the **Scanning configuration** page, For **Scan type** choose **Basic scanning**.
5. By default all of your repositories are set for **Manual** scanning. You can optionally configure scan on push by specifying **Scan on push filters**. You can set scan on push for all repositories

or individual repositories. For more information, see Filters to choose which repositories are scanned in Amazon ECR.

Manually scanning an image for OS vulnerabilities in Amazon ECR

If your repositories aren't configured to **scan on push**, you can manually start image scans. An image can be scanned once per 24 hours. The 24 hours includes the initial scan on push, if configured, and any manual scans.

For troubleshooting details for some common issues when scanning images, see Troubleshooting image scanning in Amazon ECR.

Amazon Web Services Management Console

Use the following steps to start a manual image scan using the Amazon Web Services Management Console.

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region to create your repository in.
3. In the navigation pane, choose **Repositories**.
4. On the **Repositories** page, choose the repository that contains the image to scan.
5. On the **Images** page, select the image to scan and then choose **Scan**.

Amazon CLI

- start-image-scan (Amazon CLI)

  The following example uses an image tag.

  ```
  aws ecr start-image-scan --repository-name name --image-id imageTag=tag_name --
  region us-east-2
  ```

  The following example uses an image digest.

  ```
  aws ecr start-image-scan --repository-name name --image-id imageDigest=sha256_hash
   --region us-east-2
  ```

Amazon Tools for Windows PowerShell

- Get-ECRImageScanFinding (Amazon Tools for Windows PowerShell)

  The following example uses an image tag.

  ```
  Start-ECRImageScan -RepositoryName name -ImageId_ImageTag tag_name -Region us-
  east-2 -Force
  ```

  The following example uses an image digest.

  ```
  Start-ECRImageScan -RepositoryName name -ImageId_ImageDigest sha256_hash -
  Region us-east-2 -Force
  ```

Retrieving the findings for basic scans in Amazon ECR

You can retrieve the scan findings for the last completed basic image scan. The software vulnerabilities that were discovered are listed by severity based on the Common Vulnerabilities and Exposures (CVEs) database.

For troubleshooting details for some common issues when scanning images, see Troubleshooting image scanning in Amazon ECR.

Amazon Web Services Management Console

Use the following steps to retrieve image scan findings using the Amazon Web Services Management Console.

To retrieve image scan findings

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region to create your repository in.
3. In the navigation pane, choose **Repositories**.
4. On the **Repositories** page, choose the repository that contains the image to retrieve the scan findings for.
5. On the **Images** page, under the **Vulnerabilities** column, select **Details** for the image to retrieve the scan findings for.

Amazon CLI

Use the following Amazon CLI command to retrieve image scan findings using the Amazon CLI. You can specify an image using the `imageTag` or `imageDigest`, both of which can be obtained using the list-images CLI command.

- describe-image-scan-findings (Amazon CLI)

  The following example uses an image tag.

  ```
  aws ecr describe-image-scan-findings --repository-name name --image-id
   imageTag=tag_name --region us-east-2
  ```

  The following example uses an image digest.

  ```
  aws ecr describe-image-scan-findings --repository-name name --image-id
   imageDigest=sha256_hash --region us-east-2
  ```

Amazon Tools for Windows PowerShell

- Get-ECRImageScanFinding (Amazon Tools for Windows PowerShell)

  The following example uses an image tag.

  ```
  Get-ECRImageScanFinding -RepositoryName name -ImageId_ImageTag tag_name -
  Region us-east-2
  ```

  The following example uses an image digest.

  ```
  Get-ECRImageScanFinding -RepositoryName name -ImageId_ImageDigest sha256_hash -
  Region us-east-2
  ```

Troubleshooting image scanning in Amazon ECR

The following are common image scan failures. You can view errors like this in the Amazon ECR console by displaying the image details or through the API or Amazon CLI by using the `DescribeImageScanFindings` API.

UnsupportedImageError

You may get an `UnsupportedImageError` error when attempting to perform a basic scan on an image that was built using an operating system that Amazon ECR doesn't support basic image scanning for. Amazon ECR supports package vulnerability scanning for major versions of Amazon Linux, Amazon Linux 2, Debian, Ubuntu, CentOS, Oracle Linux, Alpine, and RHEL Linux distributions. Once a distribution loses support from its vendor, Amazon ECR may no longer support scanning it for vulnerabilities. Amazon ECR does not support scanning images built from the Docker scratch image.

> **Important**
>
> When using enhanced scanning, Amazon Inspector supports scanning for specific operating systems and media types. For a full list, see Supported operating systems and media types in the *Amazon Inspector User Guide*.

An `UNDEFINED` severity level is returned

You may receive a scan finding that has a severity level of `UNDEFINED`. The following are the common causes for this:

- The vulnerability was not assigned a priority by the CVE source.
- The vulnerability was assigned a priority that Amazon ECR did not recognize.

To determine the severity and description of a vulnerability, you can view the CVE directly from the source.

Understanding scan status `SCAN_ELIGIBILITY_EXPIRED`

When enhanced scanning using Amazon Inspector is enabled for your private registry and you are viewing your scan vulnerabilities, you may see a scan status of `SCAN_ELIGIBILITY_EXPIRED`. The following are the most common causes of this.

- When you initially turn on enhanced scanning for your private registry, Amazon Inspector only recognizes images pushed to Amazon ECR in the last 30 days, based on the image push timestamp. Older images will have the `SCAN_ELIGIBILITY_EXPIRED` scan status. If you'd like these images to be scanned by Amazon Inspector you should push them again to your repository.

- If the **ECR re-scan duration** is changed in the Amazon Inspector console and that time elapses, the scan status of the image is changed to `inactive` with a reason code of `expired`, and all associated findings for the image are scheduled to be closed. This results in the Amazon ECR console listing the scan status as `SCAN_ELIGIBILITY_EXPIRED`.

Sync an upstream registry with an Amazon ECR private registry

Using pull through cache rules, you can sync the contents of an upstream registry with your Amazon ECR private registry.

Amazon ECR currently supports creating pull through cache rules for the following upstream registries.

- Docker Hub, Microsoft Azure Container Registry, , GitHub Container Registry, and GitLab Container Registry (Requires authentication)
- Amazon ECR Public, the Kubernetes container image registry, and Quay (Doesn't require authentication)

For GitLab Container Registry, Amazon ECR supports pull through cache only with GitLab software-as-a-service offering, GitLab.com.

For the upstream registries that require authentication, you must store your credentials in an Amazon Secrets Manager secret. The Amazon ECR console makes it easy for you to create the Secrets Manager secret for each of the authenticated upstream registries. For more information about creating a Secrets Manager secret using the Secrets Manager console, see Storing your upstream repository credentials in an Amazon Secrets Manager secret.

After you've created a pull through cache rule for the upstream registry, simply pull an image from that upstream registry using your Amazon ECR private registry URI. Amazon ECR then creates a repository and caches that image in your private registry. On your subsequent pull requests of the cached image with a given tag, Amazon ECR checks the upstream registry to see if there is a new version of the image with that specific tag and attempts to update the image in your private registry at least once every 24 hours.

Repository creation templates

Amazon ECR has added support for repository creation templates, currently in preview, which gives you the control to specify initial configurations for new repositories created by Amazon ECR on your behalf using pull through cache rules. Each template contains a repository namespace prefix which is used to match new repositories to a specific template. Templates can specify the

configuration for all repository settings including resource-based access policies, tag immutability, encryption, and lifecycle policies. The settings in a repository creation template are only applied during repository creation and don't have any effect on existing repositories or repositories created using any other method. For more information, see Templates to control repositories created during a pull through cache action.

Considerations for using pull through cache rules

Consider the following when using Amazon ECR pull through cache rules.

- Creating pull through cache rules isn't supported in the following Regions.
  - China (Beijing) (`cn-north-1`)
  - China (Ningxia) (`cn-northwest-1`)
  - Amazon GovCloud (US-East) (`us-gov-east-1`)
  - Amazon GovCloud (US-West) (`us-gov-west-1`)
- Amazon Lambda doesn't support pulling container images from Amazon ECR using a pull through cache rule.
- When pulling images using pull through cache, the Amazon ECR FIPS service endpoints aren't supported the first time an image is pulled. Using the Amazon ECR FIPS service endpoints work on subsequent pulls though.
- When a cached image is pulled through the Amazon ECR private registry URI, the image pulls are initiated by Amazon IP addresses. This ensures that the image pull doesn't count against any pull rate quotas implemented by the upstream registry.
- When a cached image is pulled through the Amazon ECR private registry URI, Amazon ECR checks the upstream repository at least once every 24 hours to verify whether the cached image is the latest version. If there is a newer image in the upstream registry, Amazon ECR attempts to update the cached image. This timer is based off the last pull of the cached image.
- If Amazon ECR is unable to update the image from the upstream registry for any reason and the image is pulled, the last cached image will still be pulled.
- When creating the Secrets Manager secret that contains the upstream registry credentials, the secret name must use the `ecr-pullthroughcache/` prefix. The secret must also be in the same account and Region that the pull through cache rule is created in.
- When a multi-architecture image is pulled using a pull through cache rule, the manifest list and each image referenced in the manifest list are pulled to the Amazon ECR repository. If you

only want to pull a specific architecture, you can pull the image using the image digest or tag associated with the architecture rather than the tag associated with the manifest list.

- Amazon ECR uses a service-linked IAM role, which provides the permissions needed for Amazon ECR to create the repository, retrieve the Secrets Manager secret value for authentication, and push the cached image on your behalf. The service-linked IAM role is created automatically when a pull through cache rule is created. For more information, see Amazon ECR service-linked role for pull through cache.
- By default, the IAM principal pulling the cached image has the permissions granted to them through their IAM policy. You may use the Amazon ECR private registry permissions policy to further scope the permissions of an IAM entity. For more information, see Using registry permissions.
- Amazon ECR repositories created using the pull through cache workflow are treated like any other Amazon ECR repository. All repository features, such as replication and image scanning are supported.
- When Amazon ECR creates a new repository on your behalf using a pull through cache action, the following default settings are applied to the repository unless there is a matching repository creation template. You can use a repository creation template to define the settings applied to repositories created by Amazon ECR on your behalf. For more information, see Templates to control repositories created during a pull through cache action.
  - Tag immutability — Turned off, tags are mutable and can be overwritten.
  - Encryption — The default `AES256` encryption is used.
  - Repository permissions — Omitted, no repository permissions policy is applied.
  - Lifecycle policy — Omitted, no lifecycle policy is applied.
  - Resource tags — Omitted, no resource tags are applied.
- Turning on image tag immutability for repositories using a pull through cache rule will prevent Amazon ECR from updating images using the same tag.
- When an image is pulled using the pull through cache rule for the first time a route to the internet may be required. There are certain circumstances in which a route to the internet is required so it's best to set up a route to avoid any failures. Thus, if you've configured Amazon ECR to use an interface VPC endpoint using Amazon PrivateLink then you need to ensure the first pull has a route to the internet. One way to do this is to create a public subnet in the same VPC, with an internet gateway, and then route all outbound traffic to the internet from their private subnet to the public subnet. Subsequent image pulls using the pull through cache rule don't require

this. For more information, see Example routing options in the *Amazon Virtual Private Cloud User Guide*.

IAM permissions required to sync an upstream registry with an Amazon ECR private registry

In addition to the Amazon ECR API permissions needed to authenticate to a private registry and to push and pull images, the following additional permissions are needed to use pull through cache rules effectively.

- `ecr:CreatePullThroughCacheRule` – Grants permission to create a pull through cache rule. This permission must be granted via an identity-based IAM policy.
- `ecr:BatchImportUpstreamImage` – Grants permission to retrieve the external image and import it to your private registry. This permission can be granted by using the private registry permissions policy, an identity-based IAM policy, or by using the resource-based repository permissions policy. For more information about using repository permissions, see Private repository policies in Amazon ECR.
- `ecr:CreateRepository` – Grants permission to create a repository in a private registry. This permission is required if the repository storing the cached images doesn't already exist. This permission can be granted by either an identity-based IAM policy or the private registry permissions policy.
- `ecr:TagResource` – Grants permission to add metadata tags to an Amazon ECR resource. This permission is only required if you are pulling an image that uses a pull through cache rule that has an associated repository creation template that is configured to add resource tags to the repository. This permission must be granted via an identity-based IAM policy.

Using registry permissions

Amazon ECR private registry permissions may be used to scope the permissions of individual IAM entities to use pull through cache. If an IAM entity has more permissions granted by an IAM policy than the registry permissions policy is granting, the IAM policy takes precedence. For example, if user has `ecr:*` permissions granted, no additional permissions are needed at the registry level.

To create a private registry permissions policy (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.

2. From the navigation bar, choose the Region to configure your private registry permissions statement in.
3. In the navigation pane, choose **Private registry**, **Registry permissions**.
4. On the **Registry permissions** page, choose **Generate statement**.
5. For each pull through cache permissions policy statement you want to create, do the following.
   a. For **Policy type**, choose **Pull through cache policy**.
   b. For **Statement id**, provide a name for the pull through cache statement policy.
   c. For **IAM entities**, specify the users, groups, or roles to include in the policy.
   d. For **Repository namespace**, select the pull through cache rule to associate the policy with.
   e. For **Repository names**, specify the repository base name to apply the rule for. For example, if you want to specify the Amazon Linux repository on Amazon ECR Public, the repository name would be `amazonlinux`.

To create a private registry permissions policy (Amazon CLI)

Use the following Amazon CLI command to specify the private registry permissions using the Amazon CLI.

1. Create a local file named `ptc-registry-policy.json` with the contents of your registry policy. The following example grants the `ecr-pull-through-cache-user` permission to create a repository and pull an image from Amazon ECR Public, which is the upstream source associated with the previously created pull through cache rule.

```
{
  "Sid": "PullThroughCacheFromReadOnlyRole",
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::111122223333:user/ecr-pull-through-cache-user"
  },
  "Action": [
    "ecr:CreateRepository",
    "ecr:BatchImportUpstreamImage"
  ],
  "Resource": "arn:aws:ecr:us-east-1:111122223333:repository/ecr-public/*"
}
```

> **Important**
>
> The `ecr-CreateRepository` permission is only required if the repository storing the cached images doesn't already exist. For example, if the repository creation action and the image pull actions are being done by separate IAM principals such as an administrator and a developer.

2. Use the put-registry-policy command to set the registry policy.

```
aws ecr put-registry-policy \
     --policy-text file://ptc-registry.policy.json
```

Next steps

Once you are ready to start using pull through cache rules, the following are the next steps.

- Create a pull through cache rule. For more information, see Creating a pull through cache rule in Amazon ECR.
- Create a repository creation template. A repository creation template gives you control to define the settings to use for new repositories created by Amazon ECR on your behalf during a pull through cache action. For more information, see Templates to control repositories created during a pull through cache action.

Creating a pull through cache rule in Amazon ECR

For each upstream registry containing images that you want to cache in your Amazon ECR private registry, you must create a pull through cache rule.

For upstream registries that require authentication, you must store the credentials in a Secrets Manager secret. You can use an existing secret or create a new secret. You can create the Secrets Manager secret in either the Amazon ECR console or the Secrets Manager console. To create a Secrets Manager secret using the Secrets Manager console instead of the Amazon ECR console, see Storing your upstream repository credentials in an Amazon Secrets Manager secret.

Prerequisites

- Verify that you have the proper IAM permissions to create pull through cache rules. For information, see IAM permissions required to sync an upstream registry with an Amazon ECR private registry.
- For upstream registries that require authentication: If you want to use an existing secret, verify that the Secrets Manager secret meets the following requirements:
  - The name of the secret begins with `ecr-pullthroughcache/`. The Amazon Web Services Management Console only displays Secrets Manager secrets with the `ecr-pullthroughcache/` prefix.
  - The account and Region that the secret is in must match the account and Region that the pull through cache rule is in.

To create a pull through cache rule (Amazon Web Services Management Console)

The following steps show how to create a pull through cache rule and a Secrets Manager secret using the Amazon ECR console. To create a secret using the Secrets Manager console, see Storing your upstream repository credentials in an Amazon Secrets Manager secret.

For Amazon ECR Public, Kubernetes container registry, or Quay

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.
2. From the navigation bar, choose the Region to configure your private registry settings in.
3. In the navigation pane, choose **Private registry**, **Pull through cache**.
4. On the **Pull through cache configuration** page, choose **Add rule**.
5. On the **Step 1: Specify a source** page, for **Registry**, choose either Amazon ECR Public, Kubernetes, or Quay from the list of upstream registries and then choose **Next**.
6. On the **Step 2: Specify a destination** page, for **Amazon ECR repository prefix**, specify the repository namespace prefix to use when caching images pulled from the source public registry and then choose **Next**. By default, a namespace is populated but a custom namespace can be specified as well.
7. On the **Step 3: Review and create** page, review the pull through cache rule configuration and then choose **Create**.

8. Repeat the previous step for each pull through cache you want to create. The pull through cache rules are created separately for each Region.

For Docker Hub

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.
2. From the navigation bar, choose the Region to configure your private registry settings in.
3. In the navigation pane, choose **Private registry**, **Pull through cache**.
4. On the **Pull through cache configuration** page, choose **Add rule**.
5. On the **Step 1: Specify a source** page, for **Registry**, choose **Docker Hub**, **Next**.
6. On the **Step 2: Configure authentication** page, for **Upstream credentials**, you must store your authentication credentials for Docker Hub in an Amazon Secrets Manager secret. You can specify an existing secret or use the Amazon ECR console to create a new secret.

   a. To use an existing secret, choose **Use an existing Amazon secret**. For **Secret name** use the drop down to select your existing secret, and then choose **Next**.

   > **Note**
   >
   > The Amazon Web Services Management Console only displays Secrets Manager secrets with names using the `ecr-pullthroughcache/` prefix. The secret must also be in the same account and Region that the pull through cache rule is created in.

   b. To create a new secret, choose **Create an Amazon secret**, do the following, then choose **Next**.

      i. For **Secret name**, specify a descriptive name for the secret. Secret names must contain 1-512 Unicode characters.

      ii. For **Docker Hub username**, specify your Docker Hub username.

      iii. For **Docker Hub access token**, specify your Docker Hub access token. For more information on creating a Docker Hub access token, see Create and manage access tokens in the Docker documentation.

7. On the **Step 3: Specify a destination** page, for **Amazon ECR repository prefix**, specify the repository namespace to use when caching images pulled from the source public registry and then choose **Next**.

By default, a namespace is populated but a custom namespace can be specified as well.

8. On the **Step 4: Review and create** page, review the pull through cache rule configuration and then choose **Create**.
9. Repeat the previous step for each pull through cache you want to create. The pull through cache rules are created separately for each Region.

For GitHub Container Registry

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.
2. From the navigation bar, choose the Region to configure your private registry settings in.
3. In the navigation pane, choose **Private registry**, **Pull through cache**.
4. On the **Pull through cache configuration** page, choose **Add rule**.
5. On the **Step 1: Specify a source** page, for **Registry**, choose **GitHub Container Registry**, **Next**.
6. On the **Step 2: Configure authentication** page, for **Upstream credentials**, you must store your authentication credentials for GitHub Container Registry in an Amazon Secrets Manager secret. You can specify an existing secret or use the Amazon ECR console to create a new secret.

   a. To use an existing secret, choose **Use an existing Amazon secret**. For **Secret name** use the drop down to select your existing secret, and then choose **Next**.

   > **Note**
   >
   > The Amazon Web Services Management Console only displays Secrets Manager secrets with names using the `ecr-pullthroughcache/` prefix. The secret must also be in the same account and Region that the pull through cache rule is created in.

   b. To create a new secret, choose **Create an Amazon secret**, do the following, then choose **Next**.

      i. For **Secret name**, specify a descriptive name for the secret. Secret names must contain 1-512 Unicode characters.

      ii. For **GitHub Container Registry username**, specify your GitHub Container Registry username.

iii. For **GitHub Container Registry access token**, specify your GitHub Container Registry access token. For more information on creating a GitHub access token, see Managing your personal access tokens in the GitHub documentation.

7. On the **Step 3: Specify a destination** page, for **Amazon ECR repository prefix**, specify the repository namespace to use when caching images pulled from the source public registry and then choose **Next**.

   By default, a namespace is populated but a custom namespace can be specified as well.

8. On the **Step 4: Review and create** page, review the pull through cache rule configuration and then choose **Create**.
9. Repeat the previous step for each pull through cache you want to create. The pull through cache rules are created separately for each Region.

For Microsoft Azure Container Registry

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.
2. From the navigation bar, choose the Region to configure your private registry settings in.
3. In the navigation pane, choose **Private registry**, **Pull through cache**.
4. On the **Pull through cache configuration** page, choose **Add rule**.
5. On the **Step 1: Specify a source** page, do the following.

   a. For **Registry**, choose **Microsoft Azure Container Registry**

   b. For **Source registry URL**, specify the name of your Microsoft Azure container registry and then choose **Next**.

   > **Important**
   >
   > You only need to specify the prefix, as the `.azurecr.io` suffix is populated on your behalf.

6. On the **Step 2: Configure authentication** page, for **Upstream credentials**, you must store your authentication credentials for Microsoft Azure Container Registry in an Amazon Secrets Manager secret. You can specify an existing secret or use the Amazon ECR console to create a new secret.

a. To use an existing secret, choose **Use an existing Amazon secret**. For **Secret name** use the drop down to select your existing secret, and then choose **Next**.

> **Note**
>
> The Amazon Web Services Management Console only displays Secrets Manager secrets with names using the `ecr-pullthroughcache/` prefix. The secret must also be in the same account and Region that the pull through cache rule is created in.

b. To create a new secret, choose **Create an Amazon secret**, do the following, then choose **Next**.

   i. For **Secret name**, specify a descriptive name for the secret. Secret names must contain 1-512 Unicode characters.

   ii. For **Microsoft Azure Container Registry username**, specify your Microsoft Azure Container Registry username.

   iii. For **Microsoft Azure Container Registry access token**, specify your Microsoft Azure Container Registry access token. For more information on creating an Microsoft Azure Container Registry access token, see Create token - portal in the Microsoft Azure documentation.

7. On the **Step 3: Specify a destination** page, for **Amazon ECR repository prefix**, specify the repository namespace to use when caching images pulled from the source public registry and then choose **Next**.

   By default, a namespace is populated but a custom namespace can be specified as well.

8. On the **Step 4: Review and create** page, review the pull through cache rule configuration and then choose **Create**.

9. Repeat the previous step for each pull through cache you want to create. The pull through cache rules are created separately for each Region.

For GitLab Container Registry

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.

2. From the navigation bar, choose the Region to configure your private registry settings in.

3. In the navigation pane, choose **Private registry**, **Pull through cache**.

4. On the **Pull through cache configuration** page, choose **Add rule**.
5. On the **Step 1: Specify a source** page, for Registry, choose GitLab Container Registry, Next.
6. On the **Step 2: Configure authentication** page, for **Upstream credentials**, you must store your authentication credentials for GitLab Container Registry in an Amazon Secrets Manager secret. You can specify an existing secret or use the Amazon ECR console to create a new secret.

   a. To use an existing secret, choose **Use an existing Amazon secret**. For **Secret name** use the drop down to select your existing secret, and then choose **Next**. For more information on creating a Secrets Manager secret using the Secrets Manager console, see Storing your upstream repository credentials in an Amazon Secrets Manager secret.

   > **Note**
   >
   > The Amazon Web Services Management Console only displays Secrets Manager secrets with names using the `ecr-pullthroughcache/` prefix. The secret must also be in the same account and Region that the pull through cache rule is created in.

   b. To create a new secret, choose **Create an Amazon secret**, do the following, then choose **Next**.

      i. For **Secret name**, specify a descriptive name for the secret. Secret names must contain 1-512 Unicode characters.

      ii. For **GitLab Container Registry username**, specify your GitLab Container Registry username.

      iii. For **GitLab Container Registry access token**, specify your GitLab Container Registry access token. For more information on creating a GitLab Container Registry access token, see Personal access tokens, Group access tokens, or Project access tokens, in the GitLab documentation.

7. On the **Step 3: Specify a destination** page, for **Amazon ECR repository prefix**, specify the repository namespace to use when caching images pulled from the source public registry and then choose **Next**.

   By default, a namespace is populated but a custom namespace can be specified as well.

8. On the **Step 4: Review and create** page, review the pull through cache rule configuration and then choose **Create**.

9. Repeat the previous step for each pull through cache you want to create. The pull through cache rules are created separately for each Region.

To create a pull through cache rule (Amazon CLI)

Use the create-pull-through-cache-rule Amazon CLI command to create a pull through cache rule for an Amazon ECR private registry. For upstream registries that require authentication, you must store the credentials in an Secrets Manager secret. To create a secret using the Secrets Manager console, see Storing your upstream repository credentials in an Amazon Secrets Manager secret.

The following examples are provided for each supported upstream registry.

For Amazon ECR Public

The following example creates a pull through cache rule for the Amazon ECR Public registry. It specifies a repository prefix of `ecr-public`, which results in each repository created using the pull through cache rule to have the naming scheme of `ecr-public/`*`upstream-repository-name`*.

```
aws ecr create-pull-through-cache-rule \
     --ecr-repository-prefix ecr-public \
     --upstream-registry-url public.ecr.aws \
     --region us-east-2
```

For Kubernetes container registry

The following example creates a pull through cache rule for the Kubernetes public registry. It specifies a repository prefix of `kubernetes`, which results in each repository created using the pull through cache rule to have the naming scheme of `kubernetes/`*`upstream-repository-name`*.

```
aws ecr create-pull-through-cache-rule \
     --ecr-repository-prefix kubernetes \
     --upstream-registry-url registry.k8s.io \
     --region us-east-2
```

For Quay

The following example creates a pull through cache rule for the Quay public registry. It specifies a repository prefix of `quay`, which results in each repository created using the pull through cache rule to have the naming scheme of `quay/`*`upstream-repository-name`*.

```
aws ecr create-pull-through-cache-rule \
     --ecr-repository-prefix quay \
     --upstream-registry-url quay.io \
     --region us-east-2
```

For Docker Hub

The following example creates a pull through cache rule for the Docker Hub registry. It specifies a repository prefix of `docker-hub`, which results in each repository created using the pull through cache rule to have the naming scheme of `docker-hub/`*`upstream-repository-name`*. You must specify the full Amazon Resource Name (ARN) of the secret containing your Docker Hub credentials.

```
aws ecr create-pull-through-cache-rule \
     --ecr-repository-prefix docker-hub \
     --upstream-registry-url registry-1.docker.io \
     --credential-arn arn:aws:secretsmanager:us-east-2:111122223333:secret:ecr-
pullthroughcache/example1234 \
     --region us-east-2
```

For GitHub Container Registry

The following example creates a pull through cache rule for the GitHub Container Registry. It specifies a repository prefix of `docker-hub`, which results in each repository created using the pull through cache rule to have the naming scheme of `github/`*`upstream-repository-name`*. You must specify the full Amazon Resource Name (ARN) of the secret containing your GitHub Container Registry credentials.

```
aws ecr create-pull-through-cache-rule \
     --ecr-repository-prefix github \
     --upstream-registry-url ghcr.io \
     --credential-arn arn:aws:secretsmanager:us-east-2:111122223333:secret:ecr-
pullthroughcache/example1234 \
     --region us-east-2
```

For Microsoft Azure Container Registry

The following example creates a pull through cache rule for the Microsoft Azure Container Registry. It specifies a repository prefix of `azure`, which results in each repository created using the pull through cache rule to have the naming scheme of `azure/`*`upstream-repository-name`*. You

must specify the full Amazon Resource Name (ARN) of the secret containing your Microsoft Azure Container Registry credentials.

```
aws ecr create-pull-through-cache-rule \
     --ecr-repository-prefix azure \
     --upstream-registry-url myregistry.azurecr.io \
     --credential-arn arn:aws:secretsmanager:us-east-2:111122223333:secret:ecr-
pullthroughcache/example1234 \
     --region us-east-2
```

For GitLab Container Registry

The following example creates a pull through cache rule for the GitLab Container Registry. It specifies a repository prefix of `gitlab`, which results in each repository created using the pull through cache rule to have the naming scheme of `gitlab/`*`upstream-repository-name`*. You must specify the full Amazon Resource Name (ARN) of the secret containing your GitLab Container Registry credentials.

```
aws ecr create-pull-through-cache-rule \
     --ecr-repository-prefix gitlab \
     --upstream-registry-url registry.gitlab.com \
     --credential-arn arn:aws:secretsmanager:us-east-2:111122223333:secret:ecr-
pullthroughcache/example1234 \
     --region us-east-2
```

Next steps

After you create your pull through cache rules, the following are the next steps:

- Create a repository creation template. A repository creation template gives you control to define the settings to use for new repositories created by Amazon ECR on your behalf during a pull through cache action. For more information, see Templates to control repositories created during a pull through cache action.
- Validate your pull through cache rules. When validating a pull through cache rule, Amazon ECR makes a network connection with the upstream registry, verifies that it can access the Secrets Manager secret containing the credentials for the upstream registry, and that authentication was successful. For more information, see Validating pull through cache rules in Amazon ECR.
- Start using your pull through cache rules. For more information, see Pulling an image with a pull through cache rule in Amazon ECR.

Templates to control repositories created during a pull through cache action

> The repository creation template feature is in preview release for Amazon ECR and is subject to change. During this public preview, only the Amazon Web Services Management Console can be used to manage your repository creation templates.

Use Amazon ECR repository creation templates to define the settings for repositories created by Amazon ECR on your behalf during a pull through cache action. The settings in a repository creation template are only applied during repository creation and don't have any effect on existing repositories or repositories created using any other method.

Repository creation templates aren't supported in the following Regions:

- China (Beijing) (`cn-north-1`)
- China (Ningxia) (`cn-northwest-1`)
- Amazon GovCloud (US-East) (`us-gov-east-1`)
- Amazon GovCloud (US-West) (`us-gov-west-1`)

How repository creation templates work

There are times when Amazon ECR needs to create a new private repository on your behalf. For example, the first time you use a pull through cache rule to retrieve the contents of an upstream repository and store it in your Amazon ECR private registry. When there isn't a repository creation template that matches your pull through cache rule, Amazon ECR uses the default settings for the new repository. These default settings include turning off tag immutability, using `AES-256` encryption, and not applying any repository or lifecycle policies.

Using a repository creation template with a prefix that matches a pull through cache rule gives you the ability to define the settings Amazon ECR applies to new repositories created through the pull through cache action. You can define the tag immutability, encryption configuration, repository permissions, lifecycle policy, and resource tags for the new repositories.

The following diagram shows the workflow that Amazon ECR uses when a repository creation template is used.

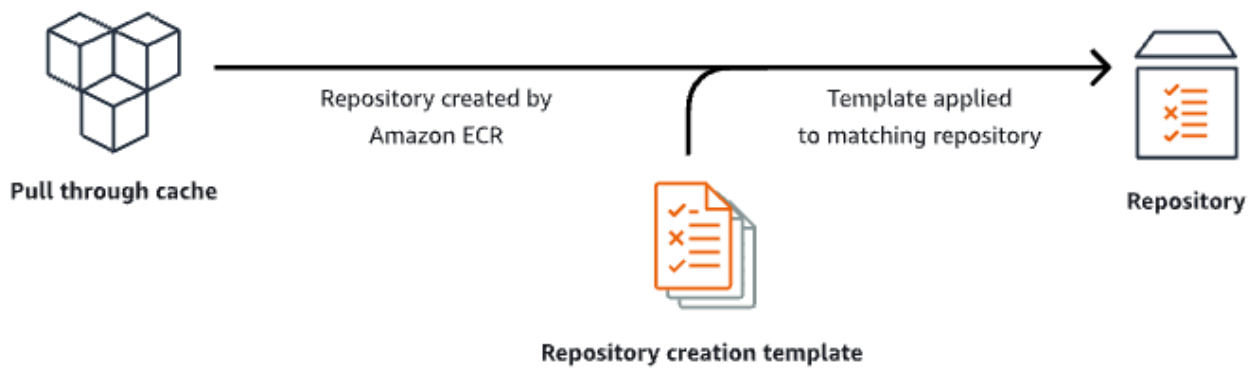

The following describes each parameter in a repository creation template in detail.

Prefix

The **Prefix** is the repository namespace prefix to associate with the template. All repositories created using this prefix will have the settings applied that are defined in this template. For example, a prefix of `prod` would apply to all repositories beginning with `prod/`. Similarly, a prefix of `prod/team` would apply to all repositories beginning with `prod/team/`.

To apply a template to all repositories in your registry that don't have an associated creation template, you can use `ROOT` as the prefix.

> **Important**
>
> There is always an assumed `/` applied to the end of the prefix. If you specify `ecr-public` as the prefix, Amazon ECR treats that as `ecr-public/`. When using a pull through cache rule, the repository prefix you specify during rule creation is what you should specify as your repository creation template prefix as well.

Description

This **template description** fis optional and is used to describe the purpose for the repository creation template.

Template version

The **repository creation template version** to use. Currently, only the `TV1` template version is supported.

Configuration version

The **repository configuration version** the template to use. Each template must include a repository configuration. The default configuration version is `CV1` and consists of the image tag mutability, repository policy, and lifecycle policy settings.

Image tag mutability

The **tag mutability** setting to use for repositories created using the template. If this parameter is omitted, the default setting of **MUTABLE** will be used which will allow image tags to be overwritten. This is the recommended setting to use for templates used for repositories created by pull through cache actions. This ensures that Amazon ECR can update the cached images when the tags are the same.

If **IMMUTABLE** is specified, all image tags within the repository will be immutable which will prevent them from being overwritten.

Encryption configuration

The **encryption configuration** to use for repositories created using the template.

If you use the **KMS** encryption type, the contents of the repository will be encrypted using server-side encryption with an Amazon Key Management Service key stored in Amazon KMS. When you use Amazon KMS to encrypt your data, you can either use the default Amazon managed Amazon KMS key for Amazon ECR, or specify your own Amazon KMS key, which you already created. For more information, see Protecting data using server-side encryption with an Amazon Key Management Service key stored in Amazon Key Management Service (SSE-KMS) in the *Amazon Simple Storage Service User Guide*.

If you use the **AES256** encryption type, Amazon ECR uses server-side encryption with Amazon S3-managed encryption keys which encrypts the images in the repository using an AES-256 encryption algorithm. For more information, see Protecting data using server-side encryption with Amazon S3-managed encryption keys (SSE-S3) in the *Amazon Simple Storage Service User Guide*.

Repository permissions

The **repository policy** to apply to repositories created using the template. A repository policy uses resource-based permissions to control access to a repository. Resource-based permissions let you specify which IAM users or roles have access to a repository and what actions they can perform on it. By default, only the Amazon account that created the repository has access to a

repository. You can apply a policy document that grant or deny additional permissions to your repository. For more information, see Private repository policies in Amazon ECR.

Repository lifecycle policy

The **lifecycle policy** to use for repositories created using the template. A lifecycle policy provides more control over the lifecycle management of images in a private repository. A lifecycle policy contains one or more rules, where each rule defines an action for Amazon ECR. This provides a way to automate the cleaning up of your container images by expiring images based on age or count. For more information, see Automate the cleanup of images by using lifecycle policies in Amazon ECR.

Resource tags

The **resource tags** are metadata to apply to the repository to help you categorize and organize them. Each tag consists of a key and an optional value, both of which you define.

IAM permissions for creating repository creation templates

The following permissions are needed for an IAM principal to manage repository creation templates. These permission must be granted using an identity-based IAM policy.

- `ecr:CreateRepositoryCreationTemplate` – Grants permission to create a repository creation template.
- `ecr:DeleteRepositoryCreationTemplate` – Grants permission to delete a repository creation template.
- `ecr:PutLifecyclePolicy` – Grants permission to create a lifecycle policy and apply it to a repository. This permission is only required if the repository creation template includes a lifecycle policy.
- `ecr:SetRepositoryPolicy` – Grants permission to create a permissions policy for a repository. This permission is only required if the repository creation template includes a repository policy.
- `ecr:TagResource` – Grants permission to add metadata tags to a resource. This permission is only required if the repository creation template includes resource tags.

Creating a repository creation template in Amazon ECR

You can create a repository creation template to define the settings to use for repositories created by Amazon ECR on your behalf during pull through cache actions. After the repository creation template is created, all new repositories created during pull through cache actions will have the settings applied. This doesn't have any effect on any previously created repositories.

To create a repository creation template (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.
2. From the navigation bar, choose the Region to create the repository creation template in.
3. In the navigation pane, choose **Private registry**, **Repository creation templates**.
4. On the **Repository creation templates** page, choose **Create template**.
5. On the **Step 1: Define template** page, for **Template details**, choose **A specific prefix** to apply the template to a specific repository namespace prefix or choose **Any prefix in your ECR registry** to apply the template to all repositories that don't match any other template in the Region.

   a. If you choose **A specific prefix**, for **Prefix** specify the repository namespace prefix to apply the template to. There is always an assumed `/` applied to the end of the prefix. For example, a prefix of `prod` would apply to all repositories beginning with `prod/`. Similarly, a prefix of `prod/team` would apply to all repositories beginning with `prod/team/`.

   b. If you choose **Any prefix in your ECR registry**, the **Prefix** will be set to `ROOT`.
6. For **Template description**, specify an optional description for the template and then choose **Next**.
7. On the **Step 2: Add repository creation configuration** page, specify the repository setting configuration to apply to repositories created using the template.

   a. For **Image tag mutability**, choose the tag mutability setting to use. For more information, see Preventing image tags from being overwritten in Amazon ECR.

      When **Mutable** is selected, image tags can be overwritten. This is the recommended setting to use for templates used for repositories created by pull through cache actions. This ensures that Amazon ECR can update the cached images when the tags are the same.

      When **Immutable** is selected, image tags are prevented from being overwritten. After the repository is configured for immutable tags, an `ImageTagAlreadyExistsException`

error is returned if there is an attempt to push an image with a tag that is already in the repository. When tag immutability is turned on for a repository, this affects all tags and you cannot make some tags immutable while others aren't.

b. For **Encryption configuration**, choose the encryption setting to use. For more information, see Encryption at rest.

   When **AES-256** is selected, Amazon ECR uses server-side encryption with Amazon Simple Storage Service-managed encryption keys which encrypts your data at rest using an industry standard AES-256 encryption algorithm. This is offered at no additional cost.

   When **Amazon KMS** is selected, Amazon ECR uses server-side encryption with keys stored in Amazon Key Management Service (Amazon KMS). When you use Amazon KMS to encrypt your data, you can either use the default Amazon managed key, which is managed by Amazon ECR, or specify your own Amazon KMS key, which is referred to as a *customer managed key*.

   > **Note**
   >
   > The encryption settings for a repository can't be changed once the repository is created.

c. For **Repository permissions**, specify the repository permissions policy to apply to repositories created using this template. You can optionally use the drop down to select one of the JSON samples for the most common use cases. For more information, see Private repository policies in Amazon ECR.

d. For **Repository lifecycle policy**, specify the repository lifecycle policy to apply to repositories created using this template. You can optionally use the drop down to select one of the JSON samples for the most common use cases. For more information, see Automate the cleanup of images by using lifecycle policies in Amazon ECR.

e. For **Repository Amazon tags**, specify the metadata, in the form of key-value pairs, to associate with the repositories created using this template and then choose **Next**. For more information, see Tagging a private repository in Amazon ECR.

8. On the **Step 3: Review and create** page, review the settings you specified for the repository creation template. Choose the **Edit** option to make changes. Choose **Create** once you're done.

Deleting a repository creation template in Amazon ECR

You can delete a repository creation template if you are finished using it. After the repository creation template is deleted, any repositories created during a pull through cache action will have default settings applied.

To delete a repository creation template (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.
2. From the navigation bar, choose the Region the repository creation template to delete is in.
3. In the navigation pane, choose **Private registry**, **Repository creation templates**.
4. On the **Repository creation templates** page, select the repository creation template to delete.
5. From the **Actions** dropdown menu, choose **Delete**.

Validating pull through cache rules in Amazon ECR

After you create a pull through cache rule, for upstream registries that require authentication you can validate that the rule works properly. When validating a pull through cache rule, Amazon ECR makes a network connection with the upstream registry, verifies that it can access the Secrets Manager secret containing the credentials for the upstream registry, and verifies that authentication was successful.

Before you start working with your pull through cache rules, verify that you have the proper IAM permissions. For more information, see IAM permissions required to sync an upstream registry with an Amazon ECR private registry.

To validate a pull through cache rule (Amazon Web Services Management Console)

The following steps show how to validate a pull through cache rule using the Amazon ECR console.

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.
2. From the navigation bar, choose the Region containing the pull through cache rule to validate.
3. In the navigation pane, choose **Private registry**, **Pull through cache**.
4. On the **Pull through cache configuration** page, select the pull through cache rule to validate. Then, use the **Actions** drop down menu and choose **View details**.

5. On the pull through cache rule detail page, use the **Actions** drop down menu and choose **Verify authentication**. Amazon ECR will display a banner with the result.
6. Repeat these steps for each pull through cache rule you want to validate.

To validate a pull through cache rule (Amazon CLI)

The validate-pull-through-cache-rule Amazon CLI command is used to validate a pull through cache rule for an Amazon ECR private registry. The following example uses the `ecr-public` namespace prefix. Replace that value with the prefix value for the pull through cache rule to validate.

```
aws ecr validate-pull-through-cache-rule \
     --ecr-repository-prefix ecr-public \
     --region us-east-2
```

In the response, the `isValid` parameter indicates whether the validation was successful or not. If `true`, Amazon ECR was able to reach the upstream registry and authentication was successful. If `false`, there was an issue and validation failed. The `failure` parameter indicates the cause.

Pulling an image with a pull through cache rule in Amazon ECR

The following examples show the command syntax to use when pulling an image using a pull through cache rule. If you receive an error pulling an upstream image using a pull through cache rule, see Troubleshooting pull through cache issues in Amazon ECR for the most common errors and how to resolve them.

Before you start working with your pull through cache rules, verify that you have the proper IAM permissions. For more information, see IAM permissions required to sync an upstream registry with an Amazon ECR private registry.

> **Note**
>
> The following examples use the default Amazon ECR repository namespace values that the Amazon Web Services Management Console uses. Ensure that you use the Amazon ECR private repository URI that you've configured.

For Amazon ECR Public

```
docker pull aws_account_id.dkr.ecr.region.amazonaws.com/ecr-public/repository_name/
image_name:tag
```

Kubernetes container registry

```
docker pull aws_account_id.dkr.ecr.region.amazonaws.com/kubernetes/repository_name/
image_name:tag
```

Quay

```
docker pull aws_account_id.dkr.ecr.region.amazonaws.com/quay/repository_name/
image_name:tag
```

Docker Hub

For Docker Hub official images:

```
docker pull aws_account_id.dkr.ecr.region.amazonaws.com/docker-hub/
library/image_name:tag
```

> **Note**
>
> For Docker Hub official images, the `/library` prefix must be included. For all other Docker Hub repositories, you should omit the `/library` prefix.

For all other Docker Hub images:

```
docker pull aws_account_id.dkr.ecr.region.amazonaws.com/docker-hub/repository_name/
image_name:tag
```

GitHub Container Registry

```
docker pull aws_account_id.dkr.ecr.region.amazonaws.com/github/repository_name/
image_name:tag
```

Microsoft Azure Container Registry

```
docker pull aws_account_id.dkr.ecr.region.amazonaws.com/azure/repository_name/
image_name:tag
```

GitLab Container Registry

```
docker pull aws_account_id.dkr.ecr.region.amazonaws.com/gitlab/repository_name/
image_name:tag
```

Storing your upstream repository credentials in an Amazon Secrets Manager secret

When creating a pull through cache rule for an upstream repository that requires authentication, you must store the credentials in a Secrets Manager secret. There may be a cost for using an Secrets Manager secret. For more information, see Amazon Secrets Manager pricing.

The following procedures walk you through how to create an Secrets Manager secret for each supported upstream repository. You can optionally use the create pull through cache rule workflow in the Amazon ECR console to create the secret instead of creating the secret using the Secrets Manager console. For more information, see Creating a pull through cache rule in Amazon ECR.

Docker Hub

To create a Secrets Manager secret for your Docker Hub credentials (Amazon Web Services Management Console)

1. Open the Secrets Manager console at https://console.amazonaws.cn/secretsmanager/.
2. Choose **Store a new secret**.
3. On the **Choose secret type** page, do the following.
   a. For **Secret type**, choose **Other type of secret**.
   b. In **Key/value pairs**, create two rows for your Docker Hub credentials. You can store up to 65536 bytes in the secret.
      i. For the first key/value pair, specify `username` as the key and your Docker Hub username as the value.

ii. For the second key/value pair, specify `accessToken` as the key and your Docker Hub access token as the value. For more information on creating a Docker Hub access token, see Create and manage access tokens in the Docker documentation.

c. For **Encryption key**, keep the default **aws/secretsmanager** Amazon KMS key value and then choose **Next**. There is no cost for using this key. For more information, see Secret encryption and decryption in Secrets Manager in the *Amazon Secrets Manager User Guide*.

> **Important**
>
> You must use the default `aws/secretsmanager` encryption key to encrypt your secret. Amazon ECR doesn't support using a customer managed key (CMK) for this.

4. On the **Configure secret** page, do the following.

a. Enter a descriptive **Secret name** and **Description**. Secret names must contain 1-512 Unicode characters and be prefixed with `ecr-pullthroughcache/`.

> **Important**
>
> The Amazon ECR Amazon Web Services Management Console only displays Secrets Manager secrets with names using the `ecr-pullthroughcache/` prefix.

b. (Optional) In the **Tags** section, add tags to your secret. For tagging strategies, see Tag Secrets Manager secrets in the *Amazon Secrets Manager User Guide*. Don't store sensitive information in tags because they aren't encrypted.

c. (Optional) In **Resource permissions**, to add a resource policy to your secret, choose **Edit permissions**. For more information, see Attach a permissions policy to an Secrets Manager secret in the *Amazon Secrets Manager User Guide*.

d. (Optional) In **Replicate secret**, to replicate your secret to another Amazon Web Services Region, choose **Replicate secret**. You can replicate your secret now or come back and replicate it later. For more information, see Replicate a secret to other Regions in the *Amazon Secrets Manager User Guide*.

e. Choose **Next**.

5. (Optional) On the **Configure rotation** page, you can turn on automatic rotation. You can also keep rotation off for now and then turn it on later. For more information, see Rotate Secrets Manager secrets in the *Amazon Secrets Manager User Guide*. Choose **Next**.
6. On the **Review** page, review your secret details, and then choose **Store**.

   Secrets Manager returns to the list of secrets. If your new secret doesn't appear, choose the refresh button.

GitHub Container Registry

To create an Secrets Manager secret for your GitHub Container Registry credentials (Amazon Web Services Management Console)

1. Open the Secrets Manager console at https://console.amazonaws.cn/secretsmanager/.
2. Choose **Store a new secret**.
3. On the **Choose secret type** page, do the following.
   a. For **Secret type**, choose **Other type of secret**.
   b. In **Key/value pairs**, create two rows for your GitHub credentials. You can store up to 65536 bytes in the secret.
      i. For the first key/value pair, specify `username` as the key and your GitHub username as the value.
      ii. For the second key/value pair, specify `accessToken` as the key and your GitHub access token as the value. For more information on creating a GitHub access token, see Managing your personal access tokens in the GitHub documentation.
   c. For **Encryption key**, keep the default **aws/secretsmanager** Amazon KMS key value and then choose **Next**. There is no cost for using this key. For more information, see Secret encryption and decryption in Secrets Manager in the *Amazon Secrets Manager User Guide*.

   > **Important**
   > You must use the default `aws/secretsmanager` encryption key to encrypt your secret. Amazon ECR doesn't support using a customer managed key (CMK) for this.

4. On the **Configure secret** page, do the following:

   a. Enter a descriptive **Secret name** and **Description**. Secret names must contain 1-512 Unicode characters and be prefixed with `ecr-pullthroughcache/`.

   > **Important**
   >
   > The Amazon ECR Amazon Web Services Management Console only displays Secrets Manager secrets with names using the `ecr-pullthroughcache/` prefix.

   b. (Optional) In the **Tags** section, add tags to your secret. For tagging strategies, see Tag Secrets Manager secrets in the *Amazon Secrets Manager User Guide*. Don't store sensitive information in tags because they aren't encrypted.

   c. (Optional) In **Resource permissions**, to add a resource policy to your secret, choose **Edit permissions**. For more information, see Attach a permissions policy to an Secrets Manager secret in the *Amazon Secrets Manager User Guide*.

   d. (Optional) In **Replicate secret**, to replicate your secret to another Amazon Web Services Region, choose **Replicate secret**. You can replicate your secret now or come back and replicate it later. For more information, see Replicate a secret to other Regions in the *Amazon Secrets Manager User Guide*.

   e. Choose **Next**.

5. (Optional) On the **Configure rotation** page, you can turn on automatic rotation. You can also keep rotation off for now and then turn it on later. For more information, see Rotate Secrets Manager secrets in the *Amazon Secrets Manager User Guide*. Choose **Next**.

6. On the **Review** page, review your secret details, and then choose **Store**.

   Secrets Manager returns to the list of secrets. If your new secret doesn't appear, choose the refresh button.

Microsoft Azure Container Registry

To create an Secrets Manager secret for your Microsoft Azure Container Registry credentials (Amazon Web Services Management Console)

1. Open the Secrets Manager console at https://console.amazonaws.cn/secretsmanager/.
2. Choose **Store a new secret**.

3. On the **Choose secret type** page, do the following.

   a. For **Secret type**, choose **Other type of secret**.

   b. In **Key/value pairs**, create two rows for your Microsoft Azure credentials. You can store up to 65536 bytes in the secret.

      i. For the first key/value pair, specify `username` as the key and your Microsoft Azure Container Registry username as the value.

      ii. For the second key/value pair, specify `accessToken` as the key and your Microsoft Azure Container Registry access token as the value. For more information on creating an Microsoft Azure access token, see Create token - portal in the Microsoft Azure documentation.

   c. For **Encryption key**, keep the default **aws/secretsmanager** Amazon KMS key value and then choose **Next**. There is no cost for using this key. For more information, see Secret encryption and decryption in Secrets Manager in the *Amazon Secrets Manager User Guide*.

   > **Important**
   >
   > You must use the default `aws/secretsmanager` encryption key to encrypt your secret. Amazon ECR doesn't support using a customer managed key (CMK) for this.

4. On the **Configure secret** page, do the following:

   a. Enter a descriptive **Secret name** and **Description**. Secret names must contain 1-512 Unicode characters and be prefixed with `ecr-pullthroughcache/`.

   > **Important**
   >
   > The Amazon ECR Amazon Web Services Management Console only displays Secrets Manager secrets with names using the `ecr-pullthroughcache/` prefix.

   b. (Optional) In the **Tags** section, add tags to your secret. For tagging strategies, see Tag Secrets Manager secrets in the *Amazon Secrets Manager User Guide*. Don't store sensitive information in tags because they aren't encrypted.

c. (Optional) In **Resource permissions**, to add a resource policy to your secret, choose **Edit permissions**. For more information, see Attach a permissions policy to an Secrets Manager secret in the *Amazon Secrets Manager User Guide*.

d. (Optional) In **Replicate secret**, to replicate your secret to another Amazon Web Services Region, choose **Replicate secret**. You can replicate your secret now or come back and replicate it later. For more information, see Replicate a secret to other Regions in the *Amazon Secrets Manager User Guide*.

e. Choose **Next**.

5. (Optional) On the **Configure rotation** page, you can turn on automatic rotation. You can also keep rotation off for now and then turn it on later. For more information, see Rotate Secrets Manager secrets in the *Amazon Secrets Manager User Guide*. Choose **Next**.

6. On the **Review** page, review your secret details, and then choose **Store**.

   Secrets Manager returns to the list of secrets. If your new secret doesn't appear, choose the refresh button.

GitLab Container Registry

To create an Secrets Manager secret for your GitLab Container Registry credentials (Amazon Web Services Management Console)

1. Open the Secrets Manager console at https://console.amazonaws.cn/secretsmanager/.

2. Choose **Store a new secret**.

3. On the **Choose secret type** page, do the following.

   a. For **Secret type**, choose **Other type of secret**.

   b. In **Key/value pairs**, create two rows for your GitLab credentials. You can store up to 65536 bytes in the secret.

      i. For the first key/value pair, specify `username` as the key and your GitLab Container Registry username as the value.

      ii. For the second key/value pair, specify `accessToken` as the key and your GitLab Container Registry access token as the value. For more information on creating a GitLab Container Registry access token, see Personal access tokens, Group access tokens, or Project access tokens, in the GitLab documentation.

c. For **Encryption key**, keep the default **aws/secretsmanager** Amazon KMS key value and then choose **Next**. There is no cost for using this key. For more information, see Secret encryption and decryption in Secrets Manager in the *Amazon Secrets Manager User Guide*.

> **Important**
>
> You must use the default `aws/secretsmanager` encryption key to encrypt your secret. Amazon ECR doesn't support using a customer managed key (CMK) for this.

4. On the **Configure secret** page, do the following:

   a. Enter a descriptive **Secret name** and **Description**. Secret names must contain 1-512 Unicode characters and be prefixed with `ecr-pullthroughcache/`.

   > **Important**
   >
   > The Amazon ECR Amazon Web Services Management Console only displays Secrets Manager secrets with names using the `ecr-pullthroughcache/` prefix.

   b. (Optional) In the **Tags** section, add tags to your secret. For tagging strategies, see Tag Secrets Manager secrets in the *Amazon Secrets Manager User Guide*. Don't store sensitive information in tags because they aren't encrypted.

   c. (Optional) In **Resource permissions**, to add a resource policy to your secret, choose **Edit permissions**. For more information, see Attach a permissions policy to an Secrets Manager secret in the *Amazon Secrets Manager User Guide*.

   d. (Optional) In **Replicate secret**, to replicate your secret to another Amazon Web Services Region, choose **Replicate secret**. You can replicate your secret now or come back and replicate it later. For more information, see Replicate a secret to other Regions in the *Amazon Secrets Manager User Guide*.

   e. Choose **Next**.

5. (Optional) On the **Configure rotation** page, you can turn on automatic rotation. You can also keep rotation off for now and then turn it on later. For more information, see Rotate Secrets Manager secrets in the *Amazon Secrets Manager User Guide*. Choose **Next**.

6. On the **Review** page, review your secret details, and then choose **Store**.

   Secrets Manager returns to the list of secrets. If your new secret doesn't appear, choose the refresh button.

Troubleshooting pull through cache issues in Amazon ECR

When pulling an upstream image using a pull through cache rule, the following are the most common errors you may receive.

Repository does not exist

An error indicating that the repository doesn't exist is most often caused by either the repository not existing in your Amazon ECR private registry or the `ecr:CreateRepository` permission not being granted to the IAM principal pulling the upstream image. To resolve this error, you should verify that the repository URI in your pull command is correct, the required IAM permissions are granted to the IAM principal pulling the upstream image, or that the repository for the upstream image to be pushed to is created in your Amazon ECR private registry before doing the upstream image pull. For more information about the required IAM permissions, see IAM permissions required to sync an upstream registry with an Amazon ECR private registry

The following is an example of this error.

```
Error response from daemon: repository 111122223333.dkr.ecr.us-east-1.amazonaws.com/
ecr-public/amazonlinux/amazonlinux not found: name unknown: The repository with
 name 'ecr-public/amazonlinux/amazonlinux' does not exist in the registry with id
 '111122223333'
```

Requested image not found

An error indicating that the image can't be found is most often caused by either the image not existing in the upstream registry or the `ecr:BatchImportUpstreamImage` permission not being granted to the IAM principal pulling the upstream image but the repository already being created in your Amazon ECR private registry. To resolve this error, you should verify the upstream image and image tag name is correct and that it exists and the required IAM permissions are granted to the IAM principal pulling the upstream image. For more information about the required IAM permissions, see IAM permissions required to sync an upstream registry with an Amazon ECR private registry.

The following is an example of this error.

```
Error response from daemon: manifest for 111122223333.dkr.ecr.us-
east-1.amazonaws.com/ecr-public/amazonlinux/amazonlinux:latest not found: manifest
 unknown: Requested image not found
```

403 Forbidden when pulling from a Docker Hub repository

When pulling from a Docker Hub repository that is tagged as a **Docker Official Image**, you must include the `/library/` in the URI you use. For example, *`aws_account_id`*`.dkr.ecr.`*`region`*`.amazonaws.com/docker-hub`**`/library/`***`image_name:tag`*. If you omit the `/library/` for Docker Hub Offical images, a `403 Forbidden` error will be returned when you attempt to pull the image using a pull through cache rule. For more information, see Pulling an image with a pull through cache rule in Amazon ECR.

The following is an example of this error.

```
Error response from daemon: failed to resolve reference "111122223333.dkr.ecr.us-
west-2.amazonaws.com/docker-hub/amazonlinux:2023": pulling from host
 111122223333.dkr.ecr.us-west-2.amazonaws.com failed with status code [manifests
 2023]: 403 Forbidden
```

Private image replication in Amazon ECR

You can configure your Amazon ECR private registry to support the replication of your repositories. Amazon ECR supports both cross-Region and cross-account replication. For cross-account replication to occur, the destination account must configure a registry permissions policy to allow replication from the source registry to occur. For more information, see Private registry permissions in Amazon ECR.

Topics

- Considerations for private image replication
- Private image replication examples for Amazon ECR
- Configuring private image replication in Amazon ECR

Considerations for private image replication

The following should be considered when using private image replication.

- Only repository content pushed to a repository after replication is configured is replicated. Any preexisting content in a repository isn't replicated. Once replication is configured for a repository, Amazon ECR keeps the destination and source synchronized.
- The repository name will remain the same across Regions and accounts when replication has occurred. Amazon ECR doesn't support changing the repository name during replication.
- The first time you configure your private registry for replication, Amazon ECR creates a service-linked IAM role on your behalf. The service-linked IAM role grants the Amazon ECR replication service the permission it needs to create repositories and replicate images in your registry. For more information, see Using service-linked roles for Amazon ECR.
- For cross-account replication to occur, the private registry destination must grant permission to allow the source registry to replicate its images. This is done by setting a private registry permissions policy. For more information, see Private registry permissions in Amazon ECR.
- If the permission policy for a private registry are changed to remove a permission, any in-progress replications previously granted may complete.
- For cross-Region replication to occur, both the source and destination accounts must be opted-in to the Region prior to any replication actions occurring within or to that Region. For more information, see Managing Amazon Regions in the *Amazon Web Services General Reference*.

- Cross-Region replication is not supported between Amazon partitions. For example, a repository in `us-west-2` can't be replicated to `cn-north-1`. For more information about Amazon partitions, see ARN format in the *Amazon General Reference*.
- The replication configuration for a private registry may contain up to 25 unique destinations across all rules, with a maximum of 10 rules total. Each rule may contain up to 100 filters. This allows for specifying separate rules for repositories containing images used for production and testing, for example.
- The replication configuration supports filtering which repositories in a private registry are replicated by specifying a repository prefix. For an example, see Example: Configuring cross-Region replication using a repository filter.
- A replication action only occurs once per image push. For example, if you configured cross-Region replication from `us-west-2` to `us-east-1` and from `us-east-1` to `us-east-2`, an image pushed to `us-west-2` replicates to only `us-east-1`, it doesn't replicate again to `us-east-2`. This behavior applies to both cross-Region and cross-account replication.
- The majority of images replicate in less than 30 minutes, but in rare cases the replication might take longer.
- Registry replication doesn't perform any delete actions. Replicated images and repositories can be manually deleted when they are no longer being used.
- Repository policies, including IAM policies, and lifecycle policies aren't replicated and don't have any effect other than on the repository they are defined for.
- Repository settings aren't replicated. The tag immutability, image scanning, and KMS encryption settings are disabled by default on all repositories created because of a replication action. The tag immutability and image scanning setting can be changed after the repository is created. However, the setting only applies to images pushed after the setting has changed.
- If tag immutability is enabled on a repository and an image is replicated that uses the same tag as an existing image, the image is replicated but won't contain the duplicated tag. This might result in the image being untagged.

Private image replication examples for Amazon ECR

The following examples show common use cases for private image replication. If you configure replication by using the Amazon CLI, you can use the JSON examples as a starting point when you create your JSON file. If you configure replication by using the Amazon Web Services Management

Console, you will see similar JSON when you review your replication rule on the **Review and submit** page.

Example: Configuring cross-Region replication to a single destination Region

The following shows an example for configuring cross-Region replication within a single registry. This example assumes that your account ID is `111122223333` and that you're specifying this replication configuration in a Region other than `us-west-2`.

```
{
    "rules": [
        {
            "destinations": [
                {
                    "region": "us-west-2",
                    "registryId": "111122223333"
                }
            ]
        }
    ]
}
```

Example: Configuring cross-Region replication using a repository filter

The following shows an example for configuring cross-Region replication for repositories that match a prefix name value. This example assumes your account ID is `111122223333` and that you're specifying this replication configuration in a Region other than `us-west-1` and have repositories with a prefix of `prod`.

```
{
 "rules": [{
  "destinations": [{
   "region": "us-west-1",
   "registryId": "111122223333"
  }],
  "repositoryFilters": [{
   "filter": "prod",
   "filterType": "PREFIX_MATCH"
  }]
 }]
```

```
}
```

Example: Configuring cross-Region replication to multiple destination Regions

The following shows an example for configuring cross-Region replication within a single registry. This example assumes your account ID is `111122223333` and that you're specifying this replication configuration in a Region other than `us-west-1` or `us-west-2`.

```
{
    "rules": [
        {
            "destinations": [
                {
                    "region": "us-west-1",
                    "registryId": "111122223333"
                },
                {
                    "region": "us-west-2",
                    "registryId": "111122223333"
                }
            ]
        }
    ]
}
```

Example: Configuring cross-account replication

The following shows an example for configuring cross-account replication for your registry. This example configures replication to the `444455556666` account and to the `us-west-2` Region.

> **Important**
>
> For cross-account replication to occur, the destination account must configure a registry permissions policy to allow replication to occur. For more information, see Private registry permissions in Amazon ECR.

```
{
```

```
    "rules": [
        {
            "destinations": [
                {
                    "region": "us-west-2",
                    "registryId": "444455556666"
                }
            ]
        }
    ]
}
```

Example: Specifying multiple rules in a configuration

The following shows an example for configuring multiple replication rules for your registry. This example configures replication for the *111122223333* account with one rule that replicates repositories with a prefix of `prod` to the `us-west-2` Region and repositories with a prefix of `test` to the `us-east-2` Region. A replication configuration may contain up to 10 rules, with each rule specifying up to 25 destinations.

```
{
 "rules": [{
   "destinations": [{
    "region": "us-west-2",
    "registryId": "111122223333"
   }],
   "repositoryFilters": [{
    "filter": "prod",
    "filterType": "PREFIX_MATCH"
   }]
  },
  {
   "destinations": [{
    "region": "us-east-2",
    "registryId": "111122223333"
   }],
   "repositoryFilters": [{
    "filter": "test",
    "filterType": "PREFIX_MATCH"
   }]
  }
 ]
```

```
}
```

Configuring private image replication in Amazon ECR

Configure replication per Region for your private registry. You can configure cross-Region replication or cross-account replication.

For examples of how replication is commonly used, see Private image replication examples for Amazon ECR.

To configure registry replication settings (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region to configure your registry replication settings for.
3. In the navigation pane, choose **Private registry**.
4. On the **Private registry** page, on the **Replication** section, choose **Edit**.
5. On the **Replication** page, choose **Add replication rule**.
6. On the **Destination types** page, choose whether to enable cross-Region replication, cross-account replication, or both and then choose **Next**.
7. If cross-Region replication is enabled, then for **Configure destination regions**, choose one or more **Destination regions** and then choose **Next**.
8. If cross-account replication is enabled, then for **Cross-account replication**, choose the cross-account replication setting for the registry. For **Destination account**, enter the account ID for the destination account and one or more **Destination regions** to replicate to. Choose **Destination account +** to configure additional accounts as replication destinations.

   > **Important**
   >
   > For cross-account replication to occur, the destination account must configure a registry permissions policy to allow replication to occur. For more information, see Private registry permissions in Amazon ECR.

9. (Optional) On the **Add filters** page, specify one or more filters for the replication rule and then choose **Add**. Repeat this step for each filter you want to associate with the replication action.

A filter must be specified as a repository name prefix. If no filters are added, the contents of all repositories are replicated. Choose **Next** once all filters have been added.

10. On the **Review and submit** page, review the replication rule configuration and then choose **Submit rule**.

To configure registry replication settings (Amazon CLI)

1. Create a JSON file containing the replication rules to define for your registry. A replication configuration may contain up to 10 rules, with up to 25 unique destinations across all rules and 100 filters per each rule. To configure cross-Region replication within your own account, you specify your own account ID. For more examples, see Private image replication examples for Amazon ECR.

```
{
 "rules": [{
  "destinations": [{
   "region": "destination_region",
   "registryId": "destination_accountId"
  }],
  "repositoryFilters": [{
   "filter": "repository_prefix_name",
   "filterType": "PREFIX_MATCH"
  }]
 }]
}
```

2. Create a replication configuration for your registry.

```
aws ecr put-replication-configuration \
     --replication-configuration file://replication-settings.json \
     --region us-west-2
```

3. Confirm your registry settings.

```
aws ecr describe-registry \
     --region us-west-2
```

Automate the cleanup of images by using lifecycle policies in Amazon ECR

Amazon ECR lifecycle policies provide more control over the lifecycle management of images in a private repository. A lifecycle policy contains one or more rules, and each rule defines an action for Amazon ECR. Based on the expiration criteria in the lifecycle policy, images expire based on age or count within 24 hours. When Amazon ECR performs an action based on a lifecycle policy, this action is captured as an event in Amazon CloudTrail. For more information, see Logging Amazon ECR actions with Amazon CloudTrail.

How lifecycle policies work

A lifecycle policy consists of one or more rules that determine which images in a repository should be expired. When considering the use of lifecycle policies, it's important to use the lifecycle policy preview to confirm which images the lifecycle policy expires before applying it to a repository. Once a lifecycle policy is applied to a repository, you should expect that images become expired within 24 hours after they meet the expiration criteria. When Amazon ECR performs an action based on a lifecycle policy, this is captured as an event in Amazon CloudTrail. For more information, see Logging Amazon ECR actions with Amazon CloudTrail.

The following diagram shows the lifecycle policy workflow.

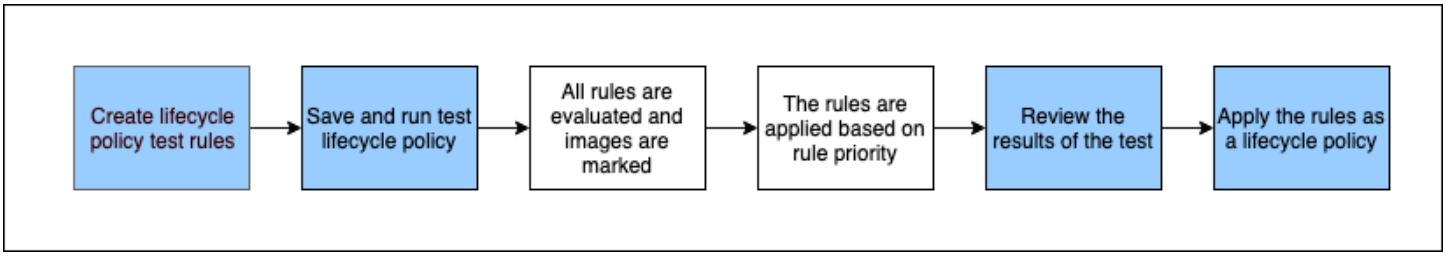

1. Create one or more test rules.
2. Save the test rules and run the preview.
3. The lifecycle policy evaluator goes through all of the rules and marks the images that each rule affects.
4. The lifecycle policy evaluator then applies the rules, based on rule priority, and displays which images in the repository are set to be expired.
5. Review the results of the test, ensuring that the images that are marked to be expired are what you intended.

6. Apply the test rules as the lifecycle policy for the repository.
7. Once the lifecycle policy is created, you should expect that images become expired within 24 hours after they meet the expiration criteria.

Lifecycle policy evaluation rules

The lifecycle policy evaluator is responsible for parsing the plaintext JSON of the lifecycle policy, evaluating all rules, and then applying those rules based on rule priority to the images in the repository. The following explains the logic of the lifecycle policy evaluator in more detail. For examples, see Examples of lifecycle policies in Amazon ECR.

- All rules are evaluated at the same time, regardless of rule priority. After all rules are evaluated, they are then applied based on rule priority.
- An image is expired by exactly one or zero rules.
- An image that matches the tagging requirements of a rule cannot be expired by a rule with a lower priority.
- Rules can never mark images that are marked by higher priority rules, but can still identify them as if they haven't been expired.
- The set of rules must contain a unique set of tag prefixes.
- Only one rule is allowed to select untagged images.
- If an image is referenced by a manifest list, it cannot be expired without the manifest list being deleted first.
- Expiration is always ordered by `pushed_at_time`, and always expires older images before newer ones.
- A lifecycle policy rule may specify either `tagPatternList` or `tagPrefixList`, but not both. However, a lifecycle policy may contain multiple rules where different rules use both pattern and prefix lists.
- The `tagPatternList` or `tagPrefixList` parameters may only used if the `tagStatus` is `tagged`.
- When using `tagPatternList`, an image is successfully matched if it matches the wildcard filter. For example, if a filter of `prod*` is applied, it would match repositories whose name begins with `prod` such as `prod`, `prod1`, or `production-team1`. Similarly, if a filter of `*prod*` is applied, it would match repositories whose name contains `prod` such as `repo-production` or `prod-team`.

> **Important**
>
> There is a maximum limit of four wildcards (*) per string. For example, `["*test*1*2*3", "test*1*2*3*"]` is valid but `["test*1*2*3*4*5*6"]` is invalid.

- When using `tagPrefixList`, an image is successfully matched if *all* of the tags in the `tagPrefixList` value are matched against any of the image's tags.
- The `countUnit` parameter is only used if `countType` is `sinceImagePushed`.
- With `countType = imageCountMoreThan`, images are sorted from youngest to oldest based on `pushed_at_time` and then all images greater than the specified count are expired.
- With `countType = sinceImagePushed`, all images whose `pushed_at_time` is older than the specified number of days based on `countNumber` are expired.

Creating a lifecycle policy preview in Amazon ECR

You can use a lifecycle policy preview to see the impact of a lifecycle policy on an image repository before you apply it. It is considered best practice to do a preview before applying a lifecycle policy to a repository.

> **Note**
>
> If you are using Amazon ECR replication to make copies of a repository across different Regions or accounts, note that a lifecycle policy can only take an action on repositories in the Region it was created in. Therefore, if you have replication turned on you may want to create a lifecycle policy in each Region and account you are replicating your repositories to.

To create a lifecycle policy preview (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region that contains the repository on which to perform a lifecycle policy preview.
3. In the navigation pane, under **Private registry**, choose **Repositories**.
4. On the **Private repositories** page, select a repository and that use the **Actions** drop down to choose **Lifecycle policies**.

5. On the lifecycle policy rules page for the repository, choose **Edit test rules**, **Create rule**.
6. Specify the following details for each test lifecycle policy rule.
   a. For **Rule priority**, type a number for the rule priority. The rule priority determines in what order the lifecycle policy rules are applied.
   b. For **Rule description**, type a description for the lifecycle policy rule.
   c. For **Image status**, choose **Tagged (wildcard matching)**, **Tagged (prefix matching)**, **Untagged**, or **Any**.
   d. If you chose **Tagged (wildcard matching)** for **Image status**, then for **Specify tags for wildcard matching**, you can specify a list of image tags with a wildcard (*) on which to take action with your lifecycle policy. For example, if your images are tagged as `prod`, `prod1`, `prod2`, and so on, you would specify `prod*` to take action on all of them. If you specify multiple tags, only the images with all specified tags are selected.

      > **Important**
      > There is a maximum limit of four wildcards (*) per string. For example, `["*test*1*2*3", "test*1*2*3*"]` is valid but `["test*1*2*3*4*5*6"]` is invalid.

   e. If you chose **Tagged (prefix matching)** for **Image status**, then for **Specify tags for prefix matching**, you can specify a list of image tags on which to take action with your lifecycle policy.
   f. For **Match criteria**, choose either **Since image pushed** or **Image count more than** and then specify a value.
   g. Choose **Save**.
7. Create additional test lifecycle policy rules by repeating steps 5–7.
8. To run the lifecycle policy preview, choose **Save and run test**.
9. Under **Image matches for test lifecycle rules**, review the impact of your lifecycle policy preview.
10. If you are satisfied with the preview results, choose **Apply as lifecycle policy** to create a lifecycle policy with the specified rules. You should expect that after applying a lifecycle policy, the affected images are expired within 24 hours.
11. If you aren't satisfied with the preview results, you may delete one or more test lifecycle rules and create one or more rules to replace them and then repeat the test.

Creating a lifecycle policy for a repository in Amazon ECR

Use a lifecycle policy to create a set of rules that expire unused repository images. After creating a lifecycle policy, the affected images are expired within 24 hours.

> **Note**
>
> If you are using Amazon ECR replication to make copies of a repository across different Regions or accounts, note that a lifecycle policy can only take an action on repositories in the Region it was created in. Therefore, if you have replication turned on you may want to create a lifecycle policy in each Region and account you are replicating your repositories to.

Prerequisite

Best practice: Create a lifecycle policy preview to verify that the images expired by your lifecycle policy rules are what you intend. For instructions, see Creating a lifecycle policy preview in Amazon ECR.

To create a lifecycle policy (Amazon Web Services Management Console)

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/repositories.
2. From the navigation bar, choose the Region that contains the repository for which to create a lifecycle policy.
3. In the navigation pane, under **Private registry**, choose **Repositories**.
4. On the **Private repositories** page, select a repository and that use the **Actions** drop down to choose **Lifecycle policies**.
5. On the lifecycle policy rules page for the repository, choose **Create rule**.
6. Enter the following details for your lifecycle policy rule.

    a. For **Rule priority**, type a number for the rule priority. The rule priority determines in what order the lifecycle policy rules are applied.

    b. For **Rule description**, type a description for the lifecycle policy rule.

    c. For **Image status**, choose **Tagged (wildcard matching)**, **Tagged (prefix matching)**, **Untagged**, or **Any**.

    d. If you chose **Tagged (wildcard matching)** for **Image status**, then for **Specify tags for wildcard matching**, you can specify a list of image tags with a wildcard (*) on which to

take action with your lifecycle policy. For example, if your images are tagged as `prod`, `prod1`, `prod2`, and so on, you would specify `prod*` to take action on all of them. If you specify multiple tags, only the images with all specified tags are selected.

> **Important**
>
> There is a maximum limit of four wildcards (*) per string. For example, `["*test*1*2*3", "test*1*2*3*"]` is valid but `["test*1*2*3*4*5*6"]` is invalid.

e. If you chose **Tagged (prefix matching)**for **Image status**, then for **Specify tags for prefix matching**, you can specify a list of image tags on which to take action with your lifecycle policy.

f. For **Match criteria**, choose either **Since image pushed** or **Image count more than** and then specify a value.

g. Choose **Save**.

7. Create additional lifecycle policy rules by repeating steps 5–7.

To create a lifecycle policy (Amazon CLI)

1. Obtain the name of the repository for which to create the lifecycle policy.

```
aws ecr describe-repositories
```

2. Create a local file named `policy.json` with the contents of the lifecycle policy. For lifecycle policy examples, see Examples of lifecycle policies in Amazon ECR.

3. Create a lifecycle policy by specifying the repository name and reference the lifecycle policy JSON file you created.

```
aws ecr put-lifecycle-policy \
      --repository-name repository-name \
      --lifecycle-policy-text file://policy.json
```

Examples of lifecycle policies in Amazon ECR

The following are example lifecycle policies showing the syntax.

To see more information about policy properties, see Lifecycle policy properties in Amazon ECR. For instructions about creating a lifecycle policy by using the Amazon CLI, see To create a lifecycle policy (Amazon CLI).

Lifecycle policy template

The contents of your lifecycle policy are evaluated before being associated with a repository. The following is the JSON syntax template for the lifecycle policy.

```
{
        "rules": [
            {
                "rulePriority": integer,
                "description": "string",
                "selection": {
                    "tagStatus": "tagged"|"untagged"|"any",
                    "tagPatternList": list<string>,
                    "tagPrefixList": list<string>,
                    "countType": "imageCountMoreThan"|"sinceImagePushed",
                    "countUnit": "string",
                    "countNumber": integer
                },
                "action": {
                    "type": "expire"
                }
            }
        ]
    }
```

Filtering on image age

The following example shows the lifecycle policy syntax for a policy that expires images with a tag starting with `prod` by using a `tagPatternList` of `prod*` that are also older than 14 days.

```
{
    "rules": [
        {
            "rulePriority": 1,
            "description": "Expire images older than 14 days",
            "selection": {
                "tagStatus": "tagged",
                "tagPatternList": ["prod*"],
```

```
                "countType": "sinceImagePushed",
                "countUnit": "days",
                "countNumber": 14
            },
            "action": {
                "type": "expire"
            }
        }
    ]
}
```

Filtering on image count

The following example shows the lifecycle policy syntax for a policy that keeps only one untagged image and expires all others.

```
{
    "rules": [
        {
            "rulePriority": 1,
            "description": "Keep only one untagged image, expire all others",
            "selection": {
                "tagStatus": "untagged",
                "countType": "imageCountMoreThan",
                "countNumber": 1
            },
            "action": {
                "type": "expire"
            }
        }
    ]
}
```

Filtering on multiple rules

The following examples use multiple rules in a lifecycle policy. An example repository and lifecycle policy are given along with an explanation of the outcome.

Example A

Repository contents:

- Image A, Taglist: ["beta-1", "prod-1"], Pushed: 10 days ago
- Image B, Taglist: ["beta-2", "prod-2"], Pushed: 9 days ago
- Image C, Taglist: ["beta-3"], Pushed: 8 days ago

Lifecycle policy text:

```
{
    "rules": [
        {
            "rulePriority": 1,
            "description": "Rule 1",
            "selection": {
                "tagStatus": "tagged",
                "tagPatternList": ["prod*"],
                "countType": "imageCountMoreThan",
                "countNumber": 1
            },
            "action": {
                "type": "expire"
            }
        },
        {
            "rulePriority": 2,
            "description": "Rule 2",
            "selection": {
                "tagStatus": "tagged",
                "tagPatternList": ["beta*"],
                "countType": "imageCountMoreThan",
                "countNumber": 1
            },
            "action": {
                "type": "expire"
            }
        }
    ]
}
```

The logic of this lifecycle policy would be:

- Rule 1 identifies images tagged with prefix `prod`. It should mark images, starting with the oldest, until there is one or fewer images remaining that match. It marks Image A for expiration.

- Rule 2 identifies images tagged with prefix `beta`. It should mark images, starting with the oldest, until there is one or fewer images remaining that match. It marks both Image A and Image B for expiration. However, Image A has already been seen by Rule 1 and if Image B were expired it would violate Rule 1 and thus is skipped.
- Result: Image A is expired.

Example B

This is the same repository as the previous example but the rule priority order is changed to illustrate the outcome.

Repository contents:

- Image A, Taglist: ["beta-1", "prod-1"], Pushed: 10 days ago
- Image B, Taglist: ["beta-2", "prod-2"], Pushed: 9 days ago
- Image C, Taglist: ["beta-3"], Pushed: 8 days ago

Lifecycle policy text:

```
{
    "rules": [
        {
            "rulePriority": 1,
            "description": "Rule 1",
            "selection": {
                "tagStatus": "tagged",
                "tagPatternList": ["beta*"],
                "countType": "imageCountMoreThan",
                "countNumber": 1
            },
            "action": {
                "type": "expire"
            }
        },
        {
            "rulePriority": 2,
            "description": "Rule 2",
            "selection": {
                "tagStatus": "tagged",
                "tagPatternList": ["prod*"],
```

```
                "countType": "imageCountMoreThan",
                "countNumber": 1
            },
            "action": {
                "type": "expire"
            }
        }
    ]
}
```

The logic of this lifecycle policy would be:

- Rule 1 identifies images tagged with prefix `beta`. It should mark images, starting with the oldest, until there is one or fewer images remaining that match. It sees all three images and would mark Image A and Image B for expiration.
- Rule 2 identifies images tagged with prefix `prod`. It should mark images, starting with the oldest, until there is one or fewer images remaining that match. It would see no images because all available images were already seen by Rule 1 and thus would mark no additional images.
- Result: Images A and B are expired.

Filtering on multiple tags in a single rule

The following examples specify the lifecycle policy syntax for multiple tag patterns in a single rule. An example repository and lifecycle policy are given along with an explanation of the outcome.

Example A

When multiple tag patterns are specified on a single rule, images must match all listed tag patterns.

Repository contents:

- Image A, Taglist: ["alpha-1"], Pushed: 12 days ago
- Image B, Taglist: ["beta-1"], Pushed: 11 days ago
- Image C, Taglist: ["alpha-2", "beta-2"], Pushed: 10 days ago
- Image D, Taglist: ["alpha-3"], Pushed: 4 days ago
- Image E, Taglist: ["beta-3"], Pushed: 3 days ago
- Image F, Taglist: ["alpha-4", "beta-4"], Pushed: 2 days ago

```
{
    "rules": [
        {
            "rulePriority": 1,
            "description": "Rule 1",
            "selection": {
                "tagStatus": "tagged",
                "tagPatternList": ["alpha*", "beta*"],
                "countType": "sinceImagePushed",
                "countNumber": 5,
                "countUnit": "days"
            },
            "action": {
                "type": "expire"
            }
        }
    ]
}
```

The logic of this lifecycle policy would be:

- Rule 1 identifies images tagged with prefix `alpha` and `beta`. It sees images C and F. It should mark images that are older than five days, which would be Image C.
- Result: Image C is expired.

Example B

The following example illustrates that tags are not exclusive.

Repository contents:

- Image A, Taglist: ["alpha-1", "beta-1", "gamma-1"], Pushed: 10 days ago
- Image B, Taglist: ["alpha-2", "beta-2"], Pushed: 9 days ago
- Image C, Taglist: ["alpha-3", "beta-3", "gamma-2"], Pushed: 8 days ago

```
{
    "rules": [
        {
            "rulePriority": 1,
            "description": "Rule 1",
```

```
            "selection": {
                "tagStatus": "tagged",
                "tagPatternList": ["alpha*", "beta*"],
                "countType": "imageCountMoreThan",
                "countNumber": 1
            },
            "action": {
                "type": "expire"
            }
        }
    ]
}
```

The logic of this lifecycle policy would be:

- Rule 1 identifies images tagged with prefix `alpha` and `beta`. It sees all images. It should mark images, starting with the oldest, until there is one or fewer images remaining that match. It marks image A and B for expiration.
- Result: Images A and B are expired.

Filtering on all images

The following lifecycle policy examples specify all images with different filters. An example repository and lifecycle policy are given along with an explanation of the outcome.

Example A

The following shows the lifecycle policy syntax for a policy that applies to all rules but keeps only one image and expires all others.

Repository contents:

- Image A, Taglist: ["alpha-1"], Pushed: 4 days ago
- Image B, Taglist: ["beta-1"], Pushed: 3 days ago
- Image C, Taglist: [], Pushed: 2 days ago
- Image D, Taglist: ["alpha-2"], Pushed: 1 day ago

```
{
    "rules": [
```

```
        {
            "rulePriority": 1,
            "description": "Rule 1",
            "selection": {
                "tagStatus": "any",
                "countType": "imageCountMoreThan",
                "countNumber": 1
            },
            "action": {
                "type": "expire"
            }
        }
    ]
}
```

The logic of this lifecycle policy would be:

- Rule 1 identifies all images. It sees images A, B, C, and D. It should expire all images other than the newest one. It marks images A, B, and C for expiration.
- Result: Images A, B, and C are expired.

Example B

The following example illustrates a lifecycle policy that combines all the rule types in a single policy.

Repository contents:

- Image A, Taglist: ["alpha-1", "beta-1"], Pushed: 4 days ago
- Image B, Taglist: [], Pushed: 3 days ago
- Image C, Taglist: ["alpha-2"], Pushed: 2 days ago
- Image D, Taglist: ["git hash"], Pushed: 1 day ago
- Image E, Taglist: [], Pushed: 1 day ago

```
{
    "rules": [
        {
            "rulePriority": 1,
            "description": "Rule 1",
```

```
            "selection": {
                "tagStatus": "tagged",
                "tagPatternList": ["alpha"],
                "countType": "imageCountMoreThan",
                "countNumber": 1
            },
            "action": {
                "type": "expire"
            }
        },
        {
            "rulePriority": 2,
            "description": "Rule 2",
            "selection": {
                "tagStatus": "untagged",
                "countType": "sinceImagePushed",
                "countUnit": "days",
                "countNumber": 1
            },
            "action": {
                "type": "expire"
            }
        },
        {
            "rulePriority": 3,
            "description": "Rule 3",
            "selection": {
                "tagStatus": "any",
                "countType": "imageCountMoreThan",
                "countNumber": 1
            },
            "action": {
                "type": "expire"
            }
        }
    ]
}
```

The logic of this lifecycle policy would be:

- Rule 1 identifies images tagged with prefix `alpha`. It identifies images A and C. It should keep the newest image and mark the rest for expiration. It marks image A for expiration.

- Rule 2 identifies untagged images. It identifies images B and E. It should mark all images older than one day for expiration. It marks image B for expiration.
- Rule 3 identifies all images. It identifies images A, B, C, D, and E. It should keep the newest image and mark the rest for expiration. However, it can't mark images A, B, C, or E because they were identified by higher priority rules. It marks image D for expiration.
- Result: Images A, B, and D are expired.

Lifecycle policy properties in Amazon ECR

Lifecycle policies have the following properties.

To see examples of lifecycle policies, see Examples of lifecycle policies in Amazon ECR. For instructions about creating a lifecycle policy by using the Amazon CLI, see To create a lifecycle policy (Amazon CLI).

Rule priority

`rulePriority`

Type: integer

Required: yes

Sets the order in which rules are applied, lowest to highest. A lifecycle policy rule with a priority of 1 is applied first, a rule with priority of 2 is next, and so on. When you add rules to a lifecycle policy, you must give them each a unique value for `rulePriority`. Values don't need to be sequential across rules in a policy. A rule with a `tagStatus` value of `any` must have the highest value for `rulePriority` and be evaluated last.

Description

`description`

Type: string

Required: no

(Optional) Describes the purpose of a rule within a lifecycle policy.

Tag status

`tagStatus`

Type: string

Required: yes

Determines whether the lifecycle policy rule that you are adding specifies a tag for an image. Acceptable options are `tagged`, `untagged`, or `any`. If you specify `any`, then all images have the rule evaluated against them. If you specify `tagged`, then you must also specify a `tagPrefixList` value. If you specify `untagged`, then you must omit `tagPrefixList`.

Tag pattern list

`tagPatternList`

Type: list[string]

Required: yes, if `tagStatus` is set to tagged and `tagPrefixList` isn't specified

When creating a lifecycle policy for tagged images, it's best practice to use a `tagPatternList` to specify the tags to expire. You specify a comma-separated list of image tag patterns that may contain wildcards (`*`) on which to take action with your lifecycle policy. For example, if your images are tagged as `prod`, `prod1`, `prod2`, and so on, you would use the tag pattern list `prod*` to specify all of them. If you specify multiple tags, only the images with all specified tags are selected.

> **Important**
>
> There is a maximum limit of four wildcards (`*`) per string. For example, `["*test*1*2*3", "test*1*2*3*"]` is valid but `["test*1*2*3*4*5*6"]` is invalid.

Tag prefix list

`tagPrefixList`

Type: list[string]

Required: yes, if `tagStatus` is set to tagged and `tagPatternList` isn't specified

Only used if you specified `"tagStatus": "tagged"` and you aren't specifying a `tagPatternList`. You must specify a comma-separated list of image tag prefixes on which to take action with your lifecycle policy. For example, if your images are tagged as `prod`, `prod1`, `prod2`, and so on, you would use the tag prefix `prod` to specify all of them. If you specify multiple tags, only the images with all specified tags are selected.

Count type

`countType`

Type: string

Required: yes

Specify a count type to apply to the images.

If `countType` is set to `imageCountMoreThan`, you also specify `countNumber` to create a rule that sets a limit on the number of images that exist in your repository. If `countType` is set to `sinceImagePushed`, you also specify `countUnit` and `countNumber` to specify a time limit on the images that exist in your repository.

Count unit

`countUnit`

Type: string

Required: yes, only if `countType` is set to `sinceImagePushed`

Specify a count unit of `days` to indicate that as the unit of time, in addition to `countNumber`, which is the number of days.

This should only be specified when `countType` is `sinceImagePushed`; an error will occur if you specify a count unit when `countType` is any other value.

Count number

`countNumber`

Type: integer

Required: yes

Specify a count number. Acceptable values are positive integers (`0` is not an accepted value).

If the `countType` used is `imageCountMoreThan`, then the value is the maximum number of images that you want to retain in your repository. If the `countType` used is `sinceImagePushed`, then the value is the maximum age limit for your images.

Action

`type`

Type: string

Required: yes

Specify an action type. The supported value is `expire`.

Security in Amazon Elastic Container Registry

Cloud security at Amazon is the highest priority. As an Amazon customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between Amazon and you. The shared responsibility model describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – Amazon is responsible for protecting the infrastructure that runs Amazon services in the Amazon Cloud. Amazon also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the Amazon compliance programs. To learn about the compliance programs that apply to Amazon ECR, see Amazon Services in Scope by Compliance Program.
- **Security in the cloud** – Your responsibility is determined by the Amazon service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Amazon ECR. The following topics show you how to configure Amazon ECR to meet your security and compliance objectives. You also learn how to use other Amazon services that help you to monitor and secure your Amazon ECR resources.

Topics

- Identity and Access Management for Amazon Elastic Container Registry
- Data protection in Amazon ECR
- Compliance validation for Amazon Elastic Container Registry
- Infrastructure Security in Amazon Elastic Container Registry
- Cross-service confused deputy prevention

Identity and Access Management for Amazon Elastic Container Registry

Amazon Identity and Access Management (IAM) is an Amazon Web Service that helps an administrator securely control access to Amazon resources. IAM administrators control who can be

authenticated (signed in) and *authorized* (have permissions) to use Amazon ECR resources. IAM is an Amazon Web Service that you can use with no additional charge.

Topics

- Audience
- Authenticating with identities
- Managing access using policies
- How Amazon Elastic Container Registry works with IAM
- Amazon Elastic Container Registry Identity-based policy examples
- Using Tag-Based Access Control
- Amazon managed policies for Amazon Elastic Container Registry
- Using service-linked roles for Amazon ECR
- Troubleshooting Amazon Elastic Container Registry Identity and Access

Audience

How you use Amazon Identity and Access Management (IAM) differs, depending on the work that you do in Amazon ECR.

Service user – If you use the Amazon ECR service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Amazon ECR features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in Amazon ECR, see Troubleshooting Amazon Elastic Container Registry Identity and Access.

Service administrator – If you're in charge of Amazon ECR resources at your company, you probably have full access to Amazon ECR. It's your job to determine which Amazon ECR features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Amazon ECR, see How Amazon Elastic Container Registry works with IAM.

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Amazon ECR. To view example Amazon ECR identity-based policies that you can use in IAM, see Amazon Elastic Container Registry Identity-based policy examples.

Authenticating with identities

Authentication is how you sign in to Amazon using your identity credentials. You must be *authenticated* (signed in to Amazon) as the Amazon Web Services account root user, as an IAM user, or by assuming an IAM role.

If you access Amazon programmatically, Amazon provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use Amazon tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see Signing Amazon API requests in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, Amazon recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see Using multi-factor authentication (MFA) in Amazon in the *IAM User Guide*.

Amazon Web Services account root user

When you create an Amazon Web Services account, you begin with one sign-in identity that has complete access to all Amazon Web Services and resources in the account. This identity is called the Amazon Web Services account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see Tasks that require root user credentials in the *IAM User Guide*.

IAM users and groups

An *IAM user* is an identity within your Amazon Web Services account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see Rotate access keys regularly for use cases that require long-term credentials in the *IAM User Guide*.

An *IAM group* is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier

to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see When to create an IAM user (instead of a role) in the *IAM User Guide*.

IAM roles

An *IAM role* is an identity within your Amazon Web Services account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. You can temporarily assume an IAM role in the Amazon Web Services Management Console by switching roles. You can assume a role by calling an Amazon CLI or Amazon API operation or by using a custom URL. For more information about methods for using roles, see Using IAM roles in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see Creating a role for a third-party Identity Provider in the *IAM User Guide*.
- **Temporary IAM user permissions** – An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- **Cross-account access** – You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some Amazon Web Services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see How IAM roles differ from resource-based policies in the *IAM User Guide*.
- **Cross-service access** – Some Amazon Web Services use features in other Amazon Web Services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
  - **Forward access sessions (FAS)** – When you use an IAM user or role to perform actions in Amazon, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of

the principal calling an Amazon Web Service, combined with the requesting Amazon Web Service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other Amazon Web Services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see Forward access sessions.

- **Service role** – A service role is an IAM role that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see Creating a role to delegate permissions to an Amazon Web Service in the *IAM User Guide*.
- **Service-linked role** – A service-linked role is a type of service role that is linked to an Amazon Web Service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your Amazon Web Services account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making Amazon CLI or Amazon API requests. This is preferable to storing access keys within the EC2 instance. To assign an Amazon role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see Using an IAM role to grant permissions to applications running on Amazon EC2 instances in the *IAM User Guide*.

To learn whether to use IAM roles or IAM users, see When to create an IAM role (instead of a user) in the *IAM User Guide*.

Managing access using policies

You control access in Amazon by creating policies and attaching them to Amazon identities or resources. A policy is an object in Amazon that, when associated with an identity or resource, defines their permissions. Amazon evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in Amazon as JSON documents. For more information about the structure and contents of JSON policy documents, see Overview of JSON policies in the *IAM User Guide*.

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the `iam:GetRole` action. A user with that policy can get role information from the Amazon Web Services Management Console, the Amazon CLI, or the Amazon API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see Creating IAM policies in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your Amazon Web Services account. Managed policies include Amazon managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see Choosing between managed policies and inline policies in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must specify a principal in a resource-based policy. Principals can include accounts, users, roles, federated users, or Amazon Web Services.

Resource-based policies are inline policies that are located in that service. You can't use Amazon managed policies from IAM in a resource-based policy.

Other policy types

Amazon supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- **Permissions boundaries** – A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the `Principal` field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see Permissions boundaries for IAM entities in the *IAM User Guide*.
- **Service control policies (SCPs)** – SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in Amazon Organizations. Amazon Organizations is a service for grouping and centrally managing multiple Amazon Web Services accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each Amazon Web Services account root user. For more information about Organizations and SCPs, see How SCPs work in the *Amazon Organizations User Guide*.
- **Session policies** – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see Session policies in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how Amazon determines whether to allow a request when multiple policy types are involved, see Policy evaluation logic in the *IAM User Guide*.

How Amazon Elastic Container Registry works with IAM

Before you use IAM to manage access to Amazon ECR, you should understand what IAM features are available to use with Amazon ECR. To get a high-level view of how Amazon ECR and other Amazon services work with IAM, see Amazon Services That Work with IAM in the *IAM User Guide*.

Topics

- Amazon ECR Identity-based policies
- Amazon ECR resource-based policies
- Authorization based on Amazon ECR tags

- Amazon ECR IAM roles

Amazon ECR Identity-based policies

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. Amazon ECR supports specific actions, resources, and condition keys. To learn about all of the elements that you use in a JSON policy, see IAM JSON Policy Elements Reference in the *IAM User Guide*.

Actions

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Action` element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated Amazon API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

Policy actions in Amazon ECR use the following prefix before the action: `ecr:`. For example, to grant someone permission to create an Amazon ECR repository with the Amazon ECR `CreateRepository` API operation, you include the `ecr:CreateRepository` action in their policy. Policy statements must include either an `Action` or `NotAction` element. Amazon ECR defines its own set of actions that describe tasks that you can perform with this service.

To specify multiple actions in a single statement, separate them with commas as follows:

```
"Action": [
      "ecr:action1",
      "ecr:action2"
```

You can specify multiple actions using wildcards (*). For example, to specify all actions that begin with the word `Describe`, include the following action:

```
"Action": "ecr:Describe*"
```

To see a list of Amazon ECR actions, see Actions, Resources, and Condition Keys for Amazon Elastic Container Registry in the *IAM User Guide*.

Resources

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Resource` JSON policy element specifies the object or objects to which the action applies. Statements must include either a `Resource` or a `NotResource` element. As a best practice, specify a resource using its Amazon Resource Name (ARN). You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

An Amazon ECR repository resource has the following ARN:

```
arn:${Partition}:ecr:${Region}:${Account}:repository/${Repository-name}
```

For more information about the format of ARNs, see Amazon Resource Names (ARNs) and Amazon Service Namespaces.

For example, to specify the `my-repo` repository in the `us-east-1` Region in your statement, use the following ARN:

```
"Resource": "arn:aws:ecr:us-east-1:123456789012:repository/my-repo"
```

To specify all repositories that belong to a specific account, use the wildcard (*):

```
"Resource": "arn:aws:ecr:us-east-1:123456789012:repository/*"
```

To specify multiple resources in a single statement, separate the ARNs with commas.

```
"Resource": [
```

```
        "resource1",
        "resource2"
```

To see a list of Amazon ECR resource types and their ARNs, see Resources Defined by Amazon Elastic Container Registry in the *IAM User Guide*. To learn with which actions you can specify the ARN of each resource, see Actions Defined by Amazon Elastic Container Registry.

Condition keys

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Condition` element (or `Condition` *block*) lets you specify conditions in which a statement is in effect. The `Condition` element is optional. You can create conditional expressions that use condition operators, such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple `Condition` elements in a statement, or multiple keys in a single `Condition` element, Amazon evaluates them using a logical `AND` operation. If you specify multiple values for a single condition key, Amazon evaluates the condition using a logical `OR` operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see IAM policy elements: variables and tags in the *IAM User Guide*.

Amazon supports global condition keys and service-specific condition keys. To see all Amazon global condition keys, see Amazon global condition context keys in the *IAM User Guide*.

Amazon ECR defines its own set of condition keys and also supports using some global condition keys. To see all Amazon global condition keys, see Amazon Global Condition Context Keys in the *IAM User Guide*.

Most Amazon ECR actions support the `aws:ResourceTag` and `ecr:ResourceTag` condition keys. For more information, see Using Tag-Based Access Control.

To see a list of Amazon ECR condition keys, see Condition Keys Defined by Amazon Elastic Container Registry in the *IAM User Guide*. To learn with which actions and resources you can use a condition key, see Actions Defined by Amazon Elastic Container Registry.

Examples

To view examples of Amazon ECR identity-based policies, see Amazon Elastic Container Registry Identity-based policy examples.

Amazon ECR resource-based policies

Resource-based policies are JSON policy documents that specify what actions a specified principal can perform on an Amazon ECR resource and under what conditions. Amazon ECR supports resource-based permissions policies for Amazon ECR repositories. Resource-based policies let you grant usage permission to other accounts on a per-resource basis. You can also use a resource-based policy to allow an Amazon service to access your Amazon ECR repositories.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. Adding a cross-account principal to a resource-based policy is only half of establishing the trust relationship. When the principal and the resource are in different Amazon accounts, you must also grant the principal entity permission to access the resource. Grant permission by attaching an identity-based policy to the entity. However, if a resource-based policy grants access to a principal in the same account, no additional identity-based policy is required. For more information, see How IAM Roles Differ from Resource-based Policies in the *IAM User Guide*.

The Amazon ECR service supports only one type of resource-based policy called a *repository policy*, which is attached to a *repository*. This policy defines which principal entities (accounts, users, roles, and federated users) can perform actions on the repository. To learn how to attach a resource-based policy to a repository, see Private repository policies in Amazon ECR.

> **Note**
>
> In an Amazon ECR repository policy, the policy element `Sid` supports additional characters and spacing not supported in IAM policies.

Examples

To view examples of Amazon ECR resource-based policies, see Private repository policy examples in Amazon ECR,

Authorization based on Amazon ECR tags

You can attach tags to Amazon ECR resources or pass tags in a request to Amazon ECR. To control access based on tags, you provide tag information in the condition element of a policy using the `ecr:ResourceTag/key-name`, `aws:RequestTag/key-name`, or `aws:TagKeys` condition keys. For more information about tagging Amazon ECR resources, see Tagging a private repository in Amazon ECR.

To view an example identity-based policy for limiting access to a resource based on the tags on that resource, see Using Tag-Based Access Control.

Amazon ECR IAM roles

An IAM role is an entity within your Amazon account that has specific permissions.

Using Temporary Credentials with Amazon ECR

You can use temporary credentials to sign in with federation, assume an IAM role, or to assume a cross-account role. You obtain temporary security credentials by calling Amazon STS API operations such as AssumeRole or GetFederationToken.

Amazon ECR supports using temporary credentials.

Service-linked roles

Service-linked roles allow Amazon services to access resources in other services to complete an action on your behalf. Service-linked roles appear in your IAM account and are owned by the service. An IAM administrator can view but not edit the permissions for service-linked roles.

Amazon ECR supports service-linked roles. For more information, see Using service-linked roles for Amazon ECR.

Amazon Elastic Container Registry Identity-based policy examples

By default, users and roles don't have permission to create or modify Amazon ECR resources. They also can't perform tasks by using the Amazon Web Services Management Console, Amazon Command Line Interface (Amazon CLI), or Amazon API. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see Creating IAM policies in the *IAM User Guide*.

For details about actions and resource types defined by Amazon ECR, including the format of the ARNs for each of the resource types, see Actions, resources, and condition keys for Amazon Elastic Container Registry in the *Service Authorization Reference*.

To learn how to create an IAM identity-based policy using these example JSON policy documents, see Creating Policies on the JSON Tab in the *IAM User Guide*.

Topics

- Policy Best Practices
- Using the Amazon ECR console
- Allow Users to View Their Own Permissions
- Accessing One Amazon ECR Repository

Policy Best Practices

Identity-based policies determine whether someone can create, access, or delete Amazon ECR resources in your account. These actions can incur costs for your Amazon Web Services account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with Amazon managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *Amazon managed policies* that grant permissions for many common use cases. They are available in your Amazon Web Services account. We recommend that you reduce permissions further by defining Amazon customer managed policies that are specific to your use cases. For more information, see Amazon managed policies or Amazon managed policies for job functions in the *IAM User Guide*.
- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see Policies and permissions in IAM in the *IAM User Guide*.
- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific Amazon Web Service, such as Amazon

CloudFormation. For more information, see IAM JSON policy elements: Condition in the *IAM User Guide*.

- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see IAM Access Analyzer policy validation in the *IAM User Guide*.
- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or a root user in your Amazon Web Services account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see Configuring MFA-protected API access in the *IAM User Guide*.

For more information about best practices in IAM, see Security best practices in IAM in the *IAM User Guide*.

Using the Amazon ECR console

To access the Amazon Elastic Container Registry console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Amazon ECR resources in your Amazon account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

To ensure that those entities can still use the Amazon ECR console, add the `AmazonEC2ContainerRegistryReadOnly` Amazon managed policy to the entities. For more information, see Adding Permissions to a User in the *IAM User Guide*:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "ecr:GetAuthorizationToken",
                "ecr:BatchCheckLayerAvailability",
                "ecr:GetDownloadUrlForLayer",
                "ecr:GetRepositoryPolicy",
                "ecr:DescribeRepositories",
```

```
                "ecr:ListImages",
                "ecr:DescribeImages",
                "ecr:BatchGetImage",
                "ecr:GetLifecyclePolicy",
                "ecr:GetLifecyclePolicyPreview",
                "ecr:ListTagsForResource",
                "ecr:DescribeImageScanFindings"
            ],
            "Resource": "*"
        }
    ]
}
```

You don't need to allow minimum console permissions for users that are making calls only to the Amazon CLI or the Amazon API. Instead, allow access to only the actions that match the API operation that you're trying to perform.

Allow Users to View Their Own Permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the Amazon CLI or Amazon API.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ViewOwnUserInfo",
            "Effect": "Allow",
            "Action": [
                "iam:GetUserPolicy",
                "iam:ListGroupsForUser",
                "iam:ListAttachedUserPolicies",
                "iam:ListUserPolicies",
                "iam:GetUser"
            ],
            "Resource": ["arn:aws-cn:iam::*:user/${aws:username}"]
        },
        {
            "Sid": "NavigateInConsole",
            "Effect": "Allow",
            "Action": [
                "iam:GetGroupPolicy",
```

```
                "iam:GetPolicyVersion",
                "iam:GetPolicy",
                "iam:ListAttachedGroupPolicies",
                "iam:ListGroupPolicies",
                "iam:ListPolicyVersions",
                "iam:ListPolicies",
                "iam:ListUsers"
            ],
            "Resource": "*"
        }
    ]
}
```

Accessing One Amazon ECR Repository

In this example, you want to grant a user in your Amazon account access to one of your Amazon ECR repositories, `my-repo`. You also want to allow the user to push, pull, and list images.

```
{
   "Version":"2012-10-17",
   "Statement":[
      {
         "Sid":"ListImagesInRepository",
         "Effect":"Allow",
         "Action":[
            "ecr:ListImages"
         ],
         "Resource":"arn:aws-cn:ecr:us-east-1:123456789012:repository/my-repo"
      },
      {
         "Sid":"GetAuthorizationToken",
         "Effect":"Allow",
         "Action":[
            "ecr:GetAuthorizationToken"
         ],
         "Resource":"*"
      },
      {
         "Sid":"ManageRepositoryContents",
         "Effect":"Allow",
         "Action":[
                "ecr:BatchCheckLayerAvailability",
                "ecr:GetDownloadUrlForLayer",
```

```
                "ecr:GetRepositoryPolicy",
                "ecr:DescribeRepositories",
                "ecr:ListImages",
                "ecr:DescribeImages",
                "ecr:BatchGetImage",
                "ecr:InitiateLayerUpload",
                "ecr:UploadLayerPart",
                "ecr:CompleteLayerUpload",
                "ecr:PutImage"
            ],
            "Resource":"arn:aws-cn:ecr:us-east-1:123456789012:repository/my-repo"
        }
    ]
}
```

Using Tag-Based Access Control

The Amazon ECR CreateRepository API action enables you to specify tags when you create the repository. For more information, see Tagging a private repository in Amazon ECR.

To enable users to tag repositories on creation, they must have permissions to use the action that creates the resource (for example, `ecr:CreateRepository`). If tags are specified in the resource-creating action, Amazon performs additional authorization on the `ecr:CreateRepository` action to verify if users have permissions to create tags.

You can use tag-based access control through IAM policies. The following are examples.

The following policy would only allow a user to create or tag a repository as `key=environment,value=dev`.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowCreateTaggedRepository",
            "Effect": "Allow",
            "Action": [
                "ecr:CreateRepository"
            ],
            "Resource": "*",
            "Condition": {
                "StringEquals": {
```

```
                    "aws:RequestTag/environment": "dev"
                }
            }
        },
        {
            "Sid": "AllowTagRepository",
            "Effect": "Allow",
            "Action": [
                "ecr:TagResource"
            ],
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "aws:RequestTag/environment": "dev"
                }
            }
        }
    ]
}
```

The following policy would allow a user access to all repositories unless they were tagged as `key=environment,value=prod`.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "ecr:*",
            "Resource": "*"
        },
        {
            "Effect": "Deny",
            "Action": "ecr:*",
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "ecr:ResourceTag/environment": "prod"
                }
            }
        }
    ]
}
```

Amazon managed policies for Amazon Elastic Container Registry

An Amazon managed policy is a standalone policy that is created and administered by Amazon. Amazon managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that Amazon managed policies might not grant least-privilege permissions for your specific use cases because they're available for all Amazon customers to use. We recommend that you reduce permissions further by defining customer managed policies that are specific to your use cases.

You cannot change the permissions defined in Amazon managed policies. If Amazon updates the permissions defined in an Amazon managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. Amazon is most likely to update an Amazon managed policy when a new Amazon Web Service is launched or new API operations become available for existing services.

For more information, see Amazon managed policies in the *IAM User Guide*.

Amazon ECR provides several managed policies that you can attach to IAM identities or to Amazon EC2 instances. These managed policies allow differing levels of control over access to Amazon ECR resources and API operations. For more information about each API operation mentioned in these policies, see Actions in the *Amazon Elastic Container Registry API Reference*.

Topics

- AmazonEC2ContainerRegistryFullAccess
- AmazonEC2ContainerRegistryPowerUser
- AmazonEC2ContainerRegistryReadOnly
- AWSECRPullThroughCache_ServiceRolePolicy
- ECRReplicationServiceRolePolicy
- Amazon ECR updates to Amazon managed policies

`AmazonEC2ContainerRegistryFullAccess`

You can attach the `AmazonEC2ContainerRegistryFullAccess` policy to your IAM identities.

You can use this managed policy as a starting point to create your own IAM policy based on your specific requirements. For example, you can create a policy specifically for providing a user or

role with full administrator access to manage the use of Amazon ECR. With the Amazon ECR Lifecycle Policies feature, you can specify the lifecycle management of images in a repository. Lifecycle policy events are reported as CloudTrail events. Amazon ECR is integrated with Amazon CloudTrail so it can display your lifecycle policy events directly in the Amazon ECR console. The `AmazonEC2ContainerRegistryFullAccess` managed IAM policy includes the `cloudtrail:LookupEvents` permission to facilitate this behavior.

Permissions details

This policy includes the following permissions:

- `ecr` – Allows principals full access to all Amazon ECR APIs.
- `cloudtrail` – Allows principals to looks up management events or Amazon CloudTrail Insights events that are captured by CloudTrail.

The `AmazonEC2ContainerRegistryFullAccess` policy is as follows.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "ecr:*",
                "cloudtrail:LookupEvents"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "iam:CreateServiceLinkedRole"
            ],
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "iam:AWSServiceName": [
                        "replication.ecr.amazonaws.com"
                    ]
                }
            }
```

```
        }
    ]
}
```

AmazonEC2ContainerRegistryPowerUser

You can attach the `AmazonEC2ContainerRegistryPowerUser` policy to your IAM identities.

This policy grants administrative permissions that allow IAM users to read and write to repositories, but doesn't allow them to delete repositories or change the policy documents that are applied to them.

Permissions details

This policy includes the following permissions:

- `ecr` – Allows principals to read and write to respositores, as well as read lifecycle policies. Principals aren't granted permission to delete repositories or change the lifecycle policies that are applied to them.

The `AmazonEC2ContainerRegistryPowerUser` policy is as follows.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "ecr:GetAuthorizationToken",
                "ecr:BatchCheckLayerAvailability",
                "ecr:GetDownloadUrlForLayer",
                "ecr:GetRepositoryPolicy",
                "ecr:DescribeRepositories",
                "ecr:ListImages",
                "ecr:DescribeImages",
                "ecr:BatchGetImage",
                "ecr:GetLifecyclePolicy",
                "ecr:GetLifecyclePolicyPreview",
                "ecr:ListTagsForResource",
                "ecr:DescribeImageScanFindings",
                "ecr:InitiateLayerUpload",
                "ecr:UploadLayerPart",
```

```
                "ecr:CompleteLayerUpload",
                "ecr:PutImage"
            ],
            "Resource": "*"
        }
    ]
}
```

AmazonEC2ContainerRegistryReadOnly

You can attach the `AmazonEC2ContainerRegistryReadOnly` policy to your IAM identities.

This policy grants read-only permissions to Amazon ECR. This includes the ability to list repositories and images within the repositories. It also includes the ability to pull images from Amazon ECR with the Docker CLI.

Permissions details

This policy includes the following permissions:

- `ecr` – Allows principals to read repositories and their respective lifecycle policies.

The `AmazonEC2ContainerRegistryReadOnly` policy is as follows.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "ecr:GetAuthorizationToken",
                "ecr:BatchCheckLayerAvailability",
                "ecr:GetDownloadUrlForLayer",
                "ecr:GetRepositoryPolicy",
                "ecr:DescribeRepositories",
                "ecr:ListImages",
                "ecr:DescribeImages",
                "ecr:BatchGetImage",
                "ecr:GetLifecyclePolicy",
                "ecr:GetLifecyclePolicyPreview",
                "ecr:ListTagsForResource",
                "ecr:DescribeImageScanFindings"
```

```
            ],
            "Resource": "*"
        }
    ]
}
```

AWSECRPullThroughCache_ServiceRolePolicy

You can't attach the `AWSECRPullThroughCache_ServiceRolePolicy` managed IAM policy to your IAM entities. This policy is attached to a service-linked role that allows Amazon ECR to push images to your repositories through the pull through cache workflow. For more information, see Amazon ECR service-linked role for pull through cache.

ECRReplicationServiceRolePolicy

You can't attach the `ECRReplicationServiceRolePolicy` managed IAM policy to your IAM entities. This policy is attached to a service-linked role that allows Amazon ECR to perform actions on your behalf. For more information, see Using service-linked roles for Amazon ECR.

Amazon ECR updates to Amazon managed policies

View details about updates to Amazon managed policies for Amazon ECR since the time that this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the Amazon ECR Document history page.

Change	Description	Date
AWSECRPullThroughCache_ServiceRolePolicy – Update to an existing policy	Amazon ECR added new permissions to the `AWSECRPullThroughCache_ServiceRolePolicy` policy. These permissions allow Amazon ECR to retrieve the encrypted contents of a Secrets Manager secret. This is required when using a pull through cache rule to cache images from	November 15, 2023

Change	Description	Date
	an upstream registry that requires authentication.	
AWSECRPullThroughCache_ServiceRolePolicy – New policy	Amazon ECR added a new policy. This policy is associated with the `AWSServiceRoleForECRPullThroughCache` service-linked role for the pull through cache feature.	November 29, 2021
ECRReplicationServiceRolePolicy – New policy	Amazon ECR added a new policy. This policy is associated with the `AWSServiceRoleForECRReplication` service-linked role for the replication feature.	December 4, 2020
AmazonEC2ContainerRegistryFullAccess – Update to an existing policy	Amazon ECR added new permissions to the `AmazonEC2ContainerRegistryFullAccess` policy. These permissions allow principals to create the Amazon ECR service-linked role.	December 4, 2020
AmazonEC2ContainerRegistryReadOnly – Update to an existing policy	Amazon ECR added new permissions to the `AmazonEC2ContainerRegistryReadOnly` policy which allow principals to read lifecycle policies, list tags, and describe the scan findings for images.	December 10, 2019

Change	Description	Date
AmazonEC2ContainerRegistryPowerUser – Update to an existing policy	Amazon ECR added new permissions to the `AmazonEC2ContainerRegistryPowerUser` policy. They allow principals to read lifecycle policies, list tags, and describe the scan findings for images.	December 10, 2019
AmazonEC2ContainerRegistryFullAccess – Update to an existing policy	Amazon ECR added new permissions to the `AmazonEC2ContainerRegistryFullAccess` policy. They allow principals to look up management events or Amazon CloudTrail Insights events that are captured by CloudTrail.	November 10, 2017
AmazonEC2ContainerRegistryReadOnly – Update to an existing policy	Amazon ECR added new permissions to the `AmazonEC2ContainerRegistryReadOnly` policy. They allow principals to describe Amazon ECR images.	October 11, 2016
AmazonEC2ContainerRegistryPowerUser – Update to an existing policy	Amazon ECR added new permissions to the `AmazonEC2ContainerRegistryPowerUser` policy. They allow principals to describe Amazon ECR images.	October 11, 2016

Change	Description	Date
AmazonEC2ContainerRegistryReadOnly – New policy	Amazon ECR added a new policy which grants grants read-only permissions to Amazon ECR. These permissions include the ability to list repositories and images within the repositories. They also include the ability to pull images from Amazon ECR with the Docker CLI.	December 21, 2015
AmazonEC2ContainerRegistryPowerUser – New policy	Amazon ECR added a new policy which grants administrative permissions that allow usrs to read and write to repositories but doesn't allow them to delete repositories or change the policy documents that are applied to them.	December 21, 2015
AmazonEC2ContainerRegistryFullAccess – New policy	Amazon ECR added a new policy. This policy grants full access to Amazon ECR.	December 21, 2015
Amazon ECR started tracking changes	Amazon ECR started tracking changes for Amazon managed policies.	June 24, 2021

Using service-linked roles for Amazon ECR

Amazon Elastic Container Registry (Amazon ECR) uses Amazon Identity and Access Management (IAM) service-linked roles to provide the permissions necessary to use the replication and pull through cache features. A service-linked role is a unique type of IAM role that is linked directly to Amazon ECR. The service-linked role is predefined by Amazon ECR. It includes all of the permissions that the service requires to support the replication and pull through cache features

for your private registry. After you configure replication or pull through cache for your registry, a service-linked role is created automatically on your behalf. For more information, see Private registry settings in Amazon ECR.

A service-linked role makes setting up replication and pull through cache with Amazon ECR easier. This is because, by using it, you don’t have to manually add all the necessary permissions. Amazon ECR defines the permissions of its service-linked roles, and unless defined otherwise, only Amazon ECR can assume its roles. The defined permissions include the trust policy and the permissions policy. The permissions policy can't be attached to any other IAM entity.

You can delete the corresponding service-linked role only after disabling either replication or pull through cache on your registry. This ensures that you don't inadvertently remove the permissions Amazon ECR requires for these features.

For information about other services that support service-linked roles, see Amazon services that work with IAM. On this linked-to page, look for the services that have **Yes** in the **Service-linked role** column. Choose a **Yes** with a link to view the relevant service-linked role documentation for that service.

Topics

- Supported Regions for Amazon ECR service-linked roles
- Amazon ECR service-linked role for replication
- Amazon ECR service-linked role for pull through cache

Supported Regions for Amazon ECR service-linked roles

Amazon ECR supports using service-linked roles in all of the Regions where the Amazon ECR service is available. For more information about Amazon ECR Region availability, see Amazon Regions and Endpoints.

Amazon ECR service-linked role for replication

Service-linked role permissions for Amazon ECR

Amazon ECR uses service-linked roles named **AWSServiceRoleForECRReplication** – Allows Amazon ECR to replicate images across multiple accounts..

The AWSServiceRoleForECRReplication service-linked role trusts the following services to assume the role:

- `replication.ecr.amazonaws.com`

The following `ECRReplicationServiceRolePolicy` role permissions policy allows Amazon ECR to use the following actions on resources:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "ecr:CreateRepository",
                "ecr:ReplicateImage"
            ],
            "Resource": "*"
        }
    ]
}
```

> **Note**
>
> The `ReplicateImage` is an internal API that Amazon ECR uses for replication and can't be called directly.

You must configure permissions to allow an IAM entity (for example a user, group, or role) to create, edit, or delete a service-linked role. For more information, see Service-Linked Role Permissions in the *IAM User Guide*.

Creating a service-linked role for Amazon ECR

You don't need to manually create the Amazon ECR service-linked role. When you configure replication settings for your registry in the Amazon Web Services Management Console, the Amazon CLI, or the Amazon API, Amazon ECR creates the service-linked role for you.

If you delete this service-linked role and need to create it again, you can use the same process to recreate the role in your account. When you configure replication settings for your registry, Amazon ECR creates the service-linked role for you again.

Editing a service-linked role for Amazon ECR

Amazon ECR doesn't allow manually editing the AWSServiceRoleForECRReplication service-linked role. After you create a service-linked role, you can't change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see Editing a service-linked role in the *IAM User Guide*.

Deleting the service-linked role for Amazon ECR

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way, you don’t have an unused entity that isn't actively monitored or maintained. However, you must remove the replication configuration for your registry in every Region before you can manually delete the service-linked role.

> **Note**
>
> If you try to delete resources while the Amazon ECR service is still using the roles, your delete action might fail. If that happens, wait for a few minutes and try again.

To delete Amazon ECR resources used by the AWSServiceRoleForECRReplication

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.
2. From the navigation bar, choose the Region your replication configuration is set on.
3. In the navigation pane, choose **Private registry**.
4. On the **Private registry** page, on the **Replication configuration** section, choose **Edit**.
5. To delete all of your replication rules, choose **Delete all**. This step requires confirmation.

To manually delete the service-linked role using IAM

Use the IAM console, the Amazon CLI, or the Amazon API to delete the **AWSServiceRoleForECRReplication** service-linked role. For more information, see Deleting a Service-Linked Role in the *IAM User Guide*.

Amazon ECR service-linked role for pull through cache

Amazon ECR uses a service-linked role named **AWSServiceRoleForECRPullThroughCache** which gives permission for Amazon ECR to perform actions on your behalf to complete pull through

cache actions. For more information about pull through cache, see Sync an upstream registry with an Amazon ECR private registry.

Service-linked role permissions for Amazon ECR

The **AWSServiceRoleForECRPullThroughCache** service-linked role trusts the following service to assume the role.

- `pullthroughcache.ecr.amazonaws.com`

Permissions details

The `AWSECRPullThroughCache_ServiceRolePolicy` permissions policy is attached to the service-linked role. This managed policy grants Amazon ECR permission to perform the following actions. For more information, see `AWSECRPullThroughCache_ServiceRolePolicy`.

- `ecr` – Allows the Amazon ECR service to push images to a private repository.
- `secretsmanager:GetSecretValue` – Allows the Amazon ECR service to retrieve the encrypted contents of an Amazon Secrets Manager secret. This is required when using a pull through cache rule to cache images from an upstream registry that requires authentication in your private registry. This permission applies only to secrets with the `ecr-pullthroughcache/` name prefix.

The `AWSECRPullThroughCache_ServiceRolePolicy` policy contains the following JSON.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ECR",
            "Effect": "Allow",
            "Action": [
                "ecr:GetAuthorizationToken",
                "ecr:BatchCheckLayerAvailability",
                "ecr:InitiateLayerUpload",
                "ecr:UploadLayerPart",
                "ecr:CompleteLayerUpload",
                "ecr:PutImage"
            ],
            "Resource": "*"
```

```
        },
        {
            "Sid": "SecretsManager",
            "Effect": "Allow",
            "Action": [
                "secretsmanager:GetSecretValue"
            ],
            "Resource": "arn:aws:secretsmanager:*:*:secret:ecr-pullthroughcache/*",
            "Condition": {
                "StringEquals": {
                    "aws:ResourceAccount": "${aws:PrincipalAccount}"
                }
            }
        }
    ]
}
```

You must configure permissions to allow an IAM entity (for example a user, group, or role) to create, edit, or delete a service-linked role. For more information, see Service-linked role permissions in the *IAM User Guide*.

Creating a service-linked role for Amazon ECR

You don't need to manually create the Amazon ECR service-linked role for pull through cache. When you create a pull through cache rule for your private registry in the Amazon Web Services Management Console, the Amazon CLI, or the Amazon API, Amazon ECR creates the service-linked role for you.

If you delete this service-linked role and need to create it again, you can use the same process to recreate the role in your account. When you create a pull through cache rule for your private registry, Amazon ECR creates the service-linked role for you again if it doesn't already exist.

Editing a service-linked role for Amazon ECR

Amazon ECR doesn't allow manually editing the **AWSServiceRoleForECRPullThroughCache** service-linked role. After the service-linked role is created, you can't change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see Editing a service-linked role in the *IAM User Guide*.

Deleting the service-linked role for Amazon ECR

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way, you don't have an unused entity that isn't actively monitored or maintained. However, you must delete the pull through cache rules for your registry in every Region before you can manually delete the service-linked role.

> **Note**
>
> If you try to delete resources while the Amazon ECR service is still using the role, your delete action might fail. If that happens, wait for a few minutes and try again.

To delete Amazon ECR resources used by the AWSServiceRoleForECRPullThroughCache service-linked role

1. Open the Amazon ECR console at https://console.amazonaws.cn/ecr/.
2. From the navigation bar, choose the Region where your pull through cache rules are created.
3. In the navigation pane, choose **Private registry**.
4. On the **Private registry** page, on the **Pull through cache configuration** section, choose **Edit**.
5. For each pull through cache rule you have created, select the rule and then choose **Delete rule**.

To manually delete the service-linked role using IAM

Use the IAM console, the Amazon CLI, or the Amazon API to delete the **AWSServiceRoleForECRPullThroughCache** service-linked role. For more information, see Deleting a Service-Linked Role in the *IAM User Guide*.

Troubleshooting Amazon Elastic Container Registry Identity and Access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon ECR and IAM.

Topics

- I Am Not Authorized to Perform an Action in Amazon ECR
- I Am Not Authorized to Perform iam:PassRole

- I want to allow people outside of my Amazon Web Services account to access my Amazon ECR resources

I Am Not Authorized to Perform an Action in Amazon ECR

If you receive an error that you're not authorized to perform an action, your policies must be updated to allow you to perform the action.

The following example error occurs when the `mateojackson` IAM user tries to use the console to view details about a fictional *`my-example-widget`* resource but doesn't have the fictional `ecr:`*`GetWidget`* permissions.

```
User: arn:aws-cn:iam::123456789012:user/mateojackson is not authorized to perform:
 ecr:GetWidget on resource: my-example-widget
```

In this case, the policy for the `mateojackson` user must be updated to allow access to the *`my-example-widget`* resource by using the `ecr:`*`GetWidget`* action.

If you need help, contact your Amazon administrator. Your administrator is the person who provided you with your sign-in credentials.

I Am Not Authorized to Perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to Amazon ECR.

Some Amazon Web Services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in Amazon ECR. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws-cn:iam::123456789012:user/marymajor is not authorized to perform:
 iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your Amazon administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my Amazon Web Services account to access my Amazon ECR resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Amazon ECR supports these features, see How Amazon Elastic Container Registry works with IAM.
- To learn how to provide access to your resources across Amazon Web Services accounts that you own, see Providing access to an IAM user in another Amazon Web Services account that you own in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party Amazon Web Services accounts, see Providing access to Amazon Web Services accounts owned by third parties in the *IAM User Guide*.
- To learn how to provide access through identity federation, see Providing access to externally authenticated users (identity federation) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see How IAM roles differ from resource-based policies in the *IAM User Guide*.

Data protection in Amazon ECR

The Amazon shared responsibility model applies to data protection in Amazon Elastic Container Service. As described in this model, Amazon is responsible for protecting the global infrastructure that runs all of the Amazon Web Services Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the Amazon Web Services that you use. For more information about data privacy, see the Data Privacy FAQ.

For data protection purposes, we recommend that you protect Amazon Web Services account credentials and set up individual users with Amazon IAM Identity Center or Amazon Identity and

Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with Amazon resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with Amazon CloudTrail.
- Use Amazon encryption solutions, along with all default security controls within Amazon Web Services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-2 validated cryptographic modules when accessing Amazon through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see Federal Information Processing Standard (FIPS) 140-2.

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Amazon ECS or other Amazon Web Services using the console, API, Amazon CLI, or Amazon SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Topics

- Encryption at rest

Encryption at rest

Amazon ECR stores images in Amazon S3 buckets that Amazon ECR manages. By default, Amazon ECR uses server-side encryption with Amazon S3-managed encryption keys which encrypts your data at rest using an AES-256 encryption algorithm. This does not require any action on your part and is offered at no additional charge. For more information, see Protecting Data Using Server-Side Encryption with Amazon S3-Managed Encryption Keys (SSE-S3) in the *Amazon Simple Storage Service User Guide*.

For more control over the encryption for your Amazon ECR repositories, you can use server-side encryption with KMS keys stored in Amazon Key Management Service (Amazon KMS). When you use Amazon KMS to encrypt your data, you can either use the default Amazon managed key, which is managed by Amazon ECR, or specify your own KMS key (referred to as a customer managed key). For more information, see Protecting Data Using Server-Side Encryption with KMS keys Stored in Amazon KMS (SSE-KMS) in the *Amazon Simple Storage Service User Guide*.

Each Amazon ECR repository has an encryption configuration, which is set when the repository is created. You can use different encryption configurations on each repository. For more information, see Creating an Amazon ECR private repository to store images.

When a repository is created with Amazon KMS encryption enabled, a KMS key is used to encrypt the contents of the repository. Moreover, Amazon ECR adds an Amazon KMS grant to the KMS key with the Amazon ECR repository as the grantee principal.

The following provides a high-level understanding of how Amazon ECR is integrated with Amazon KMS to encrypt and decrypt your repositories:

1. When creating a repository, Amazon ECR sends a DescribeKey call to Amazon KMS to validate and retrieve the Amazon Resource Name (ARN) of the KMS key specified in the encryption configuration.
2. Amazon ECR sends two CreateGrant requests to Amazon KMS to create grants on the KMS key to allow Amazon ECR to encrypt and decrypt data using the data key.
3. When pushing an image, a GenerateDataKey request is made to Amazon KMS that specifies the KMS key to use for encrypting the image layer and manifest.
4. Amazon KMS generates a new data key, encrypts it under the specified KMS key, and sends the encrypted data key to be stored with the image layer metadata and the image manifest.
5. When pulling an image, a Decrypt request is made to Amazon KMS, specifying the encrypted data key.
6. Amazon KMS decrypts the encrypted data key and sends the decrypted data key to Amazon S3.
7. The data key is used to decrypt the image layer before the image layer being pulled.
8. When a repository is deleted, Amazon ECR sends two RetireGrant requests to Amazon KMS to retire the grants created for the repository.

Considerations

The following points should be considered when using Amazon KMS encryption with Amazon ECR.

- If you create your Amazon ECR repository with KMS encryption and you do not specify a KMS key, Amazon ECR uses an Amazon managed key with the alias `aws/ecr` by default. This KMS key is created in your account the first time that you create a repository with KMS encryption enabled.
- When you use KMS encryption with your own KMS key, the key must exist in the same Region as your repository.
- The grants that Amazon ECR creates on your behalf should not be revoked. If you revoke the grant that gives Amazon ECR permission to use the Amazon KMS keys in your account, Amazon ECR cannot access this data, encrypt new images pushed to the repository, or decrypt them when they are pulled. When you revoke a grant for Amazon ECR, the change occurs immediately. To revoke access rights, you should delete the repository rather than revoking the grant. When a repository is deleted, Amazon ECR retires the grants on your behalf.
- There is a cost associated with using Amazon KMS keys. For more information, see Amazon Key Management Service pricing.

Required IAM permissions

When creating or deleting an Amazon ECR repository with server-side encryption using Amazon KMS, the permissions required depend on the specific KMS key you are using.

Required IAM permissions when using the Amazon managed key for Amazon ECR

By default, when Amazon KMS encryption is enabled for an Amazon ECR repository but no KMS key is specified, the Amazon managed key for Amazon ECR is used. When the Amazon-managed KMS key for Amazon ECR is used to encrypt a repository, any principal that has permission to create a repository can also enable Amazon KMS encryption on the repository. However, the IAM principal that deletes the repository must have the `kms:RetireGrant` permission. This enables the retirement of the grants that were added to the Amazon KMS key when the repository was created.

The following example IAM policy can be added as an inline policy to a user to ensure they have the minimum permissions needed to delete a repository that has encryption enabled. The KMS key used to encrypt the repository can be specified using the resource parameter.

```
{
    "Version": "2012-10-17",
    "Id": "ecr-kms-permissions",
    "Statement": [
```

```
        {
            "Sid": "AllowAccessToRetireTheGrantsAssociatedWithTheKey",
            "Effect": "Allow",
            "Action": [
                "kms:RetireGrant"
            ],
            "Resource": "arn:aws:kms:us-
west-2:111122223333:key/b8d9ae76-080c-4043-92EXAMPLE"
        }
    ]
}
```

Required IAM permissions when using a customer managed key

When creating a repository with Amazon KMS encryption enabled using a customer managed key, there are required permissions for both the KMS keypolicy and the IAM policy for the user or role creating the repository.

When creating your own KMS key, you can either use the default key policy Amazon KMS creates or you can specify your own. To ensure that the customer managed key remains manageable by the account owner, the key policy for the KMS key should allow all Amazon KMS actions for the root user of the account. Additional scoped permissions may be added to the key policy but at minimum the root user should be given permissions to manage the KMS key. To allow the KMS key to be used only for requests that originate in Amazon ECR, you can use the kms:ViaService condition key with the `ecr.<region>.amazonaws.com` value.

The following example key policy gives the Amazon account (root user) that owns the KMS key full access to the KMS key. For more information about this example key policy, see Allows access to the Amazon account and enables IAM policies in the *Amazon Key Management Service Developer Guide*.

```
{
    "Version": "2012-10-17",
    "Id": "ecr-key-policy",
    "Statement": [
        {
            "Sid": "EnableIAMUserPermissions",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::111122223333:root"
            },
```

```
            "Action": "kms:*",
            "Resource": "*"
        }
    ]
}
```

The IAM user, IAM role, or Amazon account creating your repositories must have the `kms:CreateGrant`, `kms:RetireGrant`, and `kms:DescribeKey` permission in addition to the necessary Amazon ECR permissions.

> **Note**
>
> The `kms:RetireGrant` permission must be added to the IAM policy of the user or role creating the repository. The `kms:CreateGrant` and `kms:DescribeKey` permissions can be added to either the key policy for the KMS key or the IAM policy of user or role creating the repository. For more information on how Amazon KMS permissions work, see Amazon KMS API permissions: Actions and resources reference in the *Amazon Key Management Service Developer Guide*.

The following example IAM policy can be added as an inline policy to a user to ensure they have the minimum permissions needed to create a repository with encryption enabled and delete the repository when they are finished with it. The Amazon KMS key used to encrypt the repository can be specified using the resource parameter.

```
{
    "Version": "2012-10-17",
    "Id": "ecr-kms-permissions",
    "Statement": [
        {
            "Sid":
 "AllowAccessToCreateAndRetireTheGrantsAssociatedWithTheKeyAsWellAsDescribeTheKey",
            "Effect": "Allow",
            "Action": [
                "kms:CreateGrant",
                "kms:RetireGrant",
                "kms:DescribeKey"
            ],
            "Resource": "arn:aws:kms:us-
west-2:111122223333:key/b8d9ae76-080c-4043-92EXAMPLE"
        }
```

```
    ]
}
```

Allow a user to list KMS keys in the console when creating a repository

When using the Amazon ECR console to create a repository, you can grant permissions to enable a user to list the customer managed KMS keys in the Region when enabling encryption for the repository. The following IAM policy example shows the permissions needed to list your KMS keys and aliases when using the console.

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

Monitoring Amazon ECR interaction with Amazon KMS

You can use Amazon CloudTrail to track the requests that Amazon ECR sends to Amazon KMS on your behalf. The log entries in the CloudTrail log contain an encryption context key to make them more easily identifiable.

Amazon ECR encryption context

An *encryption context* is a set of key–value pairs that contains arbitrary nonsecret data. When you include an encryption context in a request to encrypt data, Amazon KMS cryptographically binds the encryption context to the encrypted data. To decrypt the data, you must pass in the same encryption context.

In its GenerateDataKey and Decrypt requests to Amazon KMS, Amazon ECR uses an encryption context with two name–value pairs that identify the repository and Amazon S3 bucket being used. This is shown in the following example. The names do not vary, but combined encryption context values will be different for each value.

```
"encryptionContext": {
    "aws:s3:arn": "arn:aws:s3:::us-west-2-starport-manifest-bucket/EXAMPLE1-90ab-cdef-fedc-ba987BUCKET1/
sha256:a7766145a775d39e53a713c75b6fd6d318740e70327aaa3ed5d09e0ef33fc3df",
    "aws:ecr:arn": "arn:aws:ecr:us-west-2:111122223333:repository/repository-name"
}
```

You can use the encryption context to identify these cryptographic operation in audit records and logs, such as Amazon CloudTrail and Amazon CloudWatch Logs, and as a condition for authorization in policies and grants.

The Amazon ECR encryption context consists of two name-value pairs.

- **aws:s3:arn** – The first name–value pair identifies the bucket. The key is `aws:s3:arn`. The value is the Amazon Resource Name (ARN) of the Amazon S3 bucket.

  ```
  "aws:s3:arn": "ARN of an Amazon S3 bucket"
  ```

  For example, if the ARN of the bucket is `arn:aws:s3:::us-west-2-starport-manifest-bucket/EXAMPLE1-90ab-cdef-fedc-ba987BUCKET1/sha256:a7766145a775d39e53a713c75b6fd6d318740e70327aaa3ed5d09e0ef33fc3df`, the encryption context would include the following pair.

  ```
  "arn:aws:s3:::us-west-2-starport-manifest-bucket/EXAMPLE1-90ab-cdef-fedc-ba987BUCKET1/sha256:a7766145a775d39e53a713c75b6fd6d318740e70327aaa3ed5d09e0ef33fc3df"
  ```

- **aws:ecr:arn** – The second name–value pair identifies the Amazon Resource Name (ARN) of the repository. The key is `aws:ecr:arn`. The value is the ARN of the repository.

  ```
  "aws:ecr:arn": "ARN of an Amazon ECR repository"
  ```

  For example, if the ARN of the repository is `arn:aws:ecr:us-west-2:111122223333:repository/repository-name`, the encryption context would include the following pair.

  ```
  "aws:ecr:arn": "arn:aws:ecr:us-west-2:111122223333:repository/repository-name"
  ```

Troubleshooting

When deleting an Amazon ECR repository with the console, if the repository is successfully deleted but Amazon ECR is unable to retire the grants added to your KMS key for your repository, you will receive the following error.

```
The repository [{repository-name}] has been deleted successfully but the grants created
 by the kmsKey [{kms_key}] failed to be retired
```

When this occurs, you can retire the Amazon KMS grants for the repository yourself.

To retire Amazon KMS grants for a repository manually

1. List the grants for the Amazon KMS key used for the repository. The `key-id` value is included in the error you receive from the console. You can also use the `list-keys` command to list both the Amazon managed keys and customer managed KMS keys in a specific Region in your account.

   ```
   aws kms list-grants \
        --key-id b8d9ae76-080c-4043-9237-c815bfc21dfc
        --region us-west-2
   ```

   The output include an `EncryptionContextSubset` with the Amazon Resource Name (ARN) of your repository. This can be used to determine which grant added to the key is the one you want to retire. The `GrantId` value will be used when retiring the grant in the next step.
2. Retire each grant for the Amazon KMS key added for the repository. Replace the value for *`GrantId`* with the ID of the grant from the output of the previous step.

   ```
   aws kms retire-grant \
        --key-id b8d9ae76-080c-4043-9237-c815bfc21dfc \
        --grant-id GrantId \
        --region us-west-2
   ```

Compliance validation for Amazon Elastic Container Registry

To learn whether an Amazon Web Service is within the scope of specific compliance programs, see Amazon Web Services in Scope by Compliance Program and choose the compliance program that you are interested in. For general information, see Amazon Web Services Compliance Programs.

You can download third-party audit reports using Amazon Artifact. For more information, see Downloading Reports in Amazon Artifact.

Your compliance responsibility when using Amazon Web Services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. Amazon provides the following resources to help with compliance:

- Security and Compliance Quick Start Guides – These deployment guides discuss architectural considerations and provide steps for deploying baseline environments on Amazon that are security and compliance focused.
- Amazon Compliance Resources – This collection of workbooks and guides might apply to your industry and location.
- Evaluating Resources with Rules in the *Amazon Config Developer Guide* – The Amazon Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- Amazon Security Hub – This Amazon Web Service provides a comprehensive view of your security state within Amazon. Security Hub uses security controls to evaluate your Amazon resources and to check your compliance against security industry standards and best practices. For a list of supported services and controls, see Security Hub controls reference.

Infrastructure Security in Amazon Elastic Container Registry

As a managed service, Amazon Elastic Container Registry is protected by Amazon global network security. For information about Amazon security services and how Amazon protects infrastructure, see Amazon Cloud Security. To design your Amazon environment using the best practices for infrastructure security, see Infrastructure Protection in *Security Pillar Amazon Well-Architected Framework*.

You use Amazon published API calls to access Amazon ECR through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the Amazon Security Token Service (Amazon STS) to generate temporary security credentials to sign requests.

You can call these API operations from any network location, but Amazon ECR does support resource-based access policies, which can include restrictions based on the source IP address. You can also use Amazon ECR policies to control access from specific Amazon Virtual Private Cloud (Amazon VPC) endpoints or specific VPCs. Effectively, this isolates network access to a given Amazon ECR resource from only the specific VPC within the Amazon network. For more information, see Amazon ECR interface VPC endpoints (Amazon PrivateLink).

Amazon ECR interface VPC endpoints (Amazon PrivateLink)

You can improve the security posture of your VPC by configuring Amazon ECR to use an interface VPC endpoint. VPC endpoints are powered by Amazon PrivateLink, a technology that enables you to privately access Amazon ECR APIs through private IP addresses. Amazon PrivateLink restricts all network traffic between your VPC and Amazon ECR to the Amazon network. You don't need an internet gateway, a NAT device, or a virtual private gateway.

For more information about Amazon PrivateLink and VPC endpoints, see VPC Endpoints in the *Amazon VPC User Guide*.

Considerations for Amazon ECR VPC endpoints

Before you configure VPC endpoints for Amazon ECR, be aware of the following considerations.

- To allow your Amazon ECS tasks hosted on Amazon EC2 instances to pull private images from Amazon ECR, ensure that you also create the interface VPC endpoints for Amazon ECS. For more information, see Interface VPC Endpoints (Amazon PrivateLink) in the *Amazon Elastic Container Service Developer Guide*.

  > **Important**
  >
  > Amazon ECS tasks hosted on Fargate don't require the Amazon ECS interface VPC endpoints.

- Amazon ECS tasks hosted on Fargate using Linux platform version `1.3.0` or earlier only require the **com.amazonaws.*region*.ecr.dkr** Amazon ECR VPC endpoint and the Amazon S3 gateway endpoint to take advantage of this feature.

- Amazon ECS tasks hosted on Fargate using Linux platform version `1.4.0` or later require both the **com.amazonaws.*region*.ecr.dkr** and **com.amazonaws.*region*.ecr.api** Amazon ECR VPC endpoints as well as the Amazon S3 gateway endpoint to take advantage of this feature.
- Amazon ECS tasks hosted on Fargate using Windows platform version `1.0.0` or later require both the **com.amazonaws.*region*.ecr.dkr** and **com.amazonaws.*region*.ecr.api** Amazon ECR VPC endpoints as well as the Amazon S3 gateway endpoint to take advantage of this feature.
- Amazon ECS tasks hosted on Fargate that pull container images from Amazon ECR can restrict access to the specific VPC their tasks use and to the VPC endpoint the service uses by adding condition keys to the task execution IAM role for the task. For more information, see Optional IAM Permissions for Fargate Tasks Pulling Amazon ECR Images over Interface Endpoints in the *Amazon Elastic Container Service Developer Guide*.
- Amazon ECS tasks hosted on Fargate that pull container images from Amazon ECR that also use the `awslogs` log driver to send log information to CloudWatch Logs require the CloudWatch Logs VPC endpoint. For more information, see Create the CloudWatch Logs endpoint.
- The security group attached to the VPC endpoint must allow incoming connections on port 443 from the private subnet of the VPC.
- VPC endpoints currently don't support cross-Region requests. Ensure that you create your VPC endpoints in the same Region where you plan to issue your API calls to Amazon ECR.
- VPC endpoints currently don't support Amazon ECR Public repositories. Consider using a pull through cache rule to host the public image in a private repository in the same Region as the VPC endpoint. For more information, see Sync an upstream registry with an Amazon ECR private registry.
- VPC endpoints only support Amazon provided DNS through Amazon Route 53. If you want to use your own DNS, you can use conditional DNS forwarding. For more information, see DHCP Options Sets in the *Amazon VPC User Guide*.
- If your containers have existing connections to Amazon S3, their connections might be briefly interrupted when you add the Amazon S3 gateway endpoint. If you want to avoid this interruption, create a new VPC that uses the Amazon S3 gateway endpoint and then migrate your Amazon ECS cluster and its containers into the new VPC.
- When an image is pulled using a pull through cache rule for the first time, if you've configured Amazon ECR to use an interface VPC endpoint using Amazon PrivateLink then you need to create a public subnet in the same VPC, with a NAT gateway, and then route all outbound traffic to the internet from their private subnet to the NAT gateway in order for the pull to work. Subsequent image pulls don't require this. For more information, see Scenario: Access the internet from a private subnet in the *Amazon Virtual Private Cloud User Guide*.

Considerations for Windows images

Images based on the Windows operating system include artifacts that are restricted by license from being distributed. By default, when you push Windows images to an Amazon ECR repository, the layers that include these artifacts are not pushed as they are considered *foreign layers*. When the artifacts are provided by Microsoft, the foreign layers are retrieved from Microsoft Azure infrastructure. For this reason, to enable your containers to pull these foreign layers from Azure additional steps are needed beyond creating the VPC endpoints.

It is possible to override this behavior when pushing Windows images to Amazon ECR by using the `--allow-nondistributable-artifacts` flag in the Docker daemon. When enabled, this flag will push the licensed layers to Amazon ECR which enables these images to be pulled from Amazon ECR via the VPC endpoint without additional access to Azure being required.

> **Important**
>
> Using the `--allow-nondistributable-artifacts` flag does not preclude your obligation to comply with the terms of the Windows container base image license; you cannot post Windows content for public or third-party redistribution. Usage within your own environment is allowed.

To enable the use of this flag for your Docker installation, you must modify the Docker daemon configuration file which, depending on your Docker installation, can typically be configured in settings or preferences menu under the **Docker Engine** section or by editing the `C:\ProgramData\docker\config\daemon.json` file directly.

The following is an example of the required configuration. Replace the value with the repository URI you are pushing images to.

```
{
    "allow-nondistributable-artifacts": [
        "111122223333.dkr.ecr.us-west-2.amazonaws.com"
    ]
}
```

After modifying the Docker daemon configuration file, you must restart the Docker daemon before attempting to push your image. Confirm the push worked by verifying that the base layer was pushed to your repository.

> **Note**
>
> The base layers for Windows images are large. The layer size will result in a longer time to push and additional storage costs in Amazon ECR for these images. For these reasons, we recommend only using this option when it is strictly required to reduce build times and ongoing storage costs. For example, the `mcr.microsoft.com/windows/servercore` image is approximately 1.7 GiB in size when compressed in Amazon ECR.

Create the VPC endpoints for Amazon ECR

To create the VPC endpoints for the Amazon ECR service, use the Creating an Interface Endpoint procedure in the *Amazon VPC User Guide*.

Amazon ECS tasks hosted on Amazon EC2 instances require both Amazon ECR endpoints and the Amazon S3 gateway endpoint.

Amazon ECS tasks hosted on Fargate using platform version `1.4.0` or later require both Amazon ECR VPC endpoints and the Amazon S3 gateway endpoints.

Amazon ECS tasks hosted on Fargate that use platform version `1.3.0` or earlier only require the **cn.com.amazonaws.cn-northwest-1.ecr.dkr** Amazon ECR VPC endpoint and the Amazon S3 gateway endpoints.

> **Note**
>
> The order that the endpoints are created in doesn't matter.

cn.com.amazonaws.cn-northwest-1.ecr.dkr

This endpoint is used for the Docker Registry APIs. Docker client commands such as `push` and `pull` use this endpoint.

When you create this endpoint, you must enable a private DNS hostname. To do this, ensure that the **Enable Private DNS Name** option is selected in the Amazon VPC console when you create the VPC endpoint.

cn.com.amazonaws.cn-northwest-1.ecr.api

This endpoint is used for calls to the Amazon ECR API. API actions such as `DescribeImages` and `CreateRepository` go to this endpoint.

When this endpoint is created, you have the option to enable a private DNS hostname. Enable this setting by selecting **Enable Private DNS Name** in the VPC console when you create the VPC endpoint. If you enable a private DNS hostname for the VPC endpoint, update your SDK or Amazon CLI to the latest version so that specifying an endpoint URL when using the SDK or Amazon CLI isn't necessary.

If you enable a private DNS hostname and are using an SDK or Amazon CLI version released before January 24, 2019, you must use the `--endpoint-url` parameter to specify the interface endpoints. The following example shows the format for the endpoint URL.

```
aws ecr create-repository --repository-name name --endpoint-url https://
api.ecr.region.amazonaws.com
```

If you don't enable a private DNS hostname for the VPC endpoint, you must use the `--endpoint-url` parameter specifying the VPC endpoint ID for the interface endpoint. The following example shows the format for the endpoint URL.

```
aws ecr create-repository --repository-name name --endpoint-url
 https://VPC_endpoint_ID.api.ecr.region.vpce.amazonaws.com
```

Create the Amazon S3 gateway endpoint

For your Amazon ECS tasks to pull private images from Amazon ECR, you must create a gateway endpoint for Amazon S3. The gateway endpoint is required because Amazon ECR uses Amazon S3 to store your image layers. When your containers download images from Amazon ECR, they must access Amazon ECR to get the image manifest and then Amazon S3 to download the actual image layers. The following is the Amazon Resource Name (ARN) of the Amazon S3 bucket containing the layers for each Docker image.

```
arn:aws:s3:::prod-region-starport-layer-bucket/*
```

Use the Creating a gateway endpoint procedure in the *Amazon VPC User Guide* to create the following Amazon S3 gateway endpoint for Amazon ECR. When creating the endpoint, be sure to select the route tables for your VPC.

com.amazonaws.*region*.s3

The Amazon S3 gateway endpoint uses an IAM policy document to limit access to the service. The **Full Access** policy can be used because any restrictions that you have put in your task IAM roles or other IAM user policies still apply on top of this policy. If you want to limit Amazon S3 bucket access to the minimum required permissions for using Amazon ECR, see Minimum Amazon S3 Bucket Permissions for Amazon ECR.

Minimum Amazon S3 Bucket Permissions for Amazon ECR

The Amazon S3 gateway endpoint uses an IAM policy document to limit access to the service. To allow only the minimum Amazon S3 bucket permissions for Amazon ECR, restrict access to the Amazon S3 bucket that Amazon ECR uses when you create the IAM policy document for the endpoint.

The following table describes the Amazon S3 bucket policy permissions needed by Amazon ECR.

Permission	Description

`arn:aws:s3:::prod-`*`region`*`-starport-layer-bucket/*`
 	
`arn:aws-cn:s3:::prod-`*`region`*`-starport-layer-bucket/*`
	Provides access to the Amazon S3 bucket containing the layers for each Docker image. Represents the Region identifier for an Amazon Region supported by Amazon ECR, such as `us-east-2` for the US East (Ohio) Region.

Example

The following example illustrates how to provide access to the Amazon S3 buckets required for Amazon ECR operations.

```
{
  "Statement": [
    {
```

```
      "Sid": "Access-to-specific-bucket-only",
      "Principal": "*",
      "Action": [
        "s3:GetObject"
      ],
      "Effect": "Allow",
      "Resource": ["arn:aws:s3:::prod-region-starport-layer-bucket/*"]
    }
  ]
}
```

Create the CloudWatch Logs endpoint

Amazon ECS tasks using the Fargate launch type that use a VPC without an internet gateway that also use the `awslogs` log driver to send log information to CloudWatch Logs require that you create the **com.amazonaws.*region*.logs** interface VPC endpoint for CloudWatch Logs. For more information, see Using CloudWatch Logs with interface VPC endpoints in the *Amazon CloudWatch Logs User Guide*.

Create an endpoint policy for your Amazon ECR VPC endpoints

A VPC endpoint policy is an IAM resource policy that you attach to an endpoint when you create or modify the endpoint. If you don't attach a policy when you create an endpoint, Amazon attaches a default policy for you that allows full access to the service. An endpoint policy doesn't override or replace user policies or service-specific policies. It's a separate policy for controlling access from the endpoint to the specified service. Endpoint policies must be written in JSON format. For more information, see Controlling Access to Services with VPC Endpoints in the *Amazon VPC User Guide*.

We recommend creating a single IAM resource policy and attaching it to both of the Amazon ECR VPC endpoints.

The following is an example of an endpoint policy for Amazon ECR. This policy enables a specific IAM role to pull images from Amazon ECR.

```
{
 "Statement": [{
  "Sid": "AllowPull",
  "Principal": {
   "AWS": "arn:aws:iam::1234567890:role/role_name"
  },
  "Action": [
```

```
    "ecr:BatchGetImage",
    "ecr:GetDownloadUrlForLayer",
                      "ecr:GetAuthorizationToken"
   ],
   "Effect": "Allow",
   "Resource": "*"
  }]
}
```

The following endpoint policy example prevents a specified repository from being deleted.

```
{
 "Statement": [{
   "Sid": "AllowAll",
   "Principal": "*",
   "Action": "*",
   "Effect": "Allow",
   "Resource": "*"
  },
  {
   "Sid": "PreventDelete",
   "Principal": "*",
   "Action": "ecr:DeleteRepository",
   "Effect": "Deny",
   "Resource": "arn:aws:ecr:region:1234567890:repository/repository_name"
  }
 ]
}
```

The following endpoint policy example combines the two previous examples into a single policy.

```
{
 "Statement": [{
   "Sid": "AllowAll",
   "Effect": "Allow",
   "Principal": "*",
   "Action": "*",
   "Resource": "*"
  },
  {
   "Sid": "PreventDelete",
   "Effect": "Deny",
   "Principal": "*",
```

```
    "Action": "ecr:DeleteRepository",
    "Resource": "arn:aws:ecr:region:1234567890:repository/repository_name"
   },
   {
    "Sid": "AllowPull",
    "Effect": "Allow",
    "Principal": {
     "AWS": "arn:aws:iam::1234567890:role/role_name"
    },
    "Action": [
     "ecr:BatchGetImage",
     "ecr:GetDownloadUrlForLayer",
                           "ecr:GetAuthorizationToken"
    ],
    "Resource": "*"
   }
  ]
 }
```

To modify the VPC endpoint policy for Amazon ECR

1. Open the Amazon VPC console at https://console.amazonaws.cn/vpc/.
2. In the navigation pane, choose **Endpoints**.
3. If you have not already created the VPC endpoints for Amazon ECR, see Create the VPC endpoints for Amazon ECR.
4. Select the Amazon ECR VPC endpoint to add a policy to, and choose the **Policy** tab in the lower half of the screen.
5. Choose **Edit Policy** and make the changes to the policy.
6. Choose **Save** to save the policy.

Shared subnets

You can't create, describe, modify, or delete VPC endpoints in subnets that are shared with you. However, you can use the VPC endpoints in subnets that are shared with you.

Cross-service confused deputy prevention

The confused deputy problem is a security issue where an entity that doesn't have permission to perform an action can coerce a more-privileged entity to perform the action. In Amazon, cross-

service impersonation can result in the confused deputy problem. Cross-service impersonation can occur when one service (the *calling service*) calls another service (the *called service*). The calling service can be manipulated to use its permissions to act on another customer's resources in a way it should not otherwise have permission to access. To prevent this, Amazon provides tools that help you protect your data for all services with service principals that have been given access to resources in your account.

We recommend using the `aws:SourceArn` or `aws:SourceAccount` global condition context keys in resource policies to limit the permissions that Amazon ECR gives another service to the resource. Use `aws:SourceArn` if you want only one resource to be associated with the cross-service access. Use `aws:SourceAccount` if you want to allow any resource in that account to be associated with the cross-service use.

The most effective way to protect against the confused deputy problem is to use the `aws:SourceArn` global condition context key with the full ARN of the resource. If you don't know the full ARN of the resource or if you are specifying multiple resources, use the `aws:SourceArn` global context condition key with wildcard characters (*) for the unknown portions of the ARN. For example, `arn:aws-cn:servicename:region:123456789012:*`.

If the `aws:SourceArn` value does not contain the account ID, such as an Amazon S3 bucket ARN, you must use both global condition context keys to limit permissions.

The value of `aws:SourceArn` must be ResourceDescription.

The following example shows how you can use the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in an Amazon ECR repository policy to allow Amazon CodeBuild access to the Amazon ECR API actions necessary for integration with that service while also preventing the confused deputy problem.

```
{
   "Version":"2012-10-17",
   "Statement":[
      {
         "Sid":"CodeBuildAccess",
         "Effect":"Allow",
         "Principal":{
            "Service":"codebuild.amazonaws.com"
         },
         "Action":[
            "ecr:BatchGetImage",
```

```
            "ecr:GetDownloadUrlForLayer"
         ],
         "Condition":{
            "ArnLike":{
               "aws:SourceArn":"arn:aws:codebuild:region:123456789012:project/project-
name"
            },
            "StringEquals":{
               "aws:SourceAccount":"123456789012"
            }
         }
      }
   ]
}
```

Amazon ECR monitoring

You can monitor your Amazon ECR API usage with Amazon CloudWatch, which collects and processes raw data from Amazon ECR into readable, near real-time metrics. These statistics are recorded for a period of two weeks so that you can access historical information and gain perspective on your API usage. Amazon ECR metric data is automatically sent to CloudWatch in one-minute periods. For more information about CloudWatch, see the Amazon CloudWatch User Guide.

Amazon ECR provides metrics based on your API usage for authorization, image push, and image pull actions.

Monitoring is an important part of maintaining the reliability, availability, and performance of Amazon ECR and your Amazon solutions. We recommend that you collect monitoring data from the resources that make up your Amazon solution so that you can more easily debug a multi-point failure if one occurs. Before you start monitoring Amazon ECR, however, you should create a monitoring plan that includes answers to the following questions:

- What are your monitoring goals?
- What resources will you monitor?
- How often will you monitor these resources?
- What monitoring tools will you use?
- Who will perform the monitoring tasks?
- Who should be notified when something goes wrong?

The next step is to establish a baseline for normal Amazon ECR performance in your environment by measuring performance at various times and under different load conditions. As you monitor Amazon ECR, store historical monitoring data so that you can compare it with new performance data, identify normal performance patterns and performance anomalies, and devise methods to address issues.

Topics

- Visualizing your service quotas and setting alarms
- Amazon ECR usage metrics
- Amazon ECR usage reports

- Amazon ECR repository metrics
- Amazon ECR events and EventBridge
- Logging Amazon ECR actions with Amazon CloudTrail

Visualizing your service quotas and setting alarms

You can use the CloudWatch console to visualize your service quotas and see how your current usage compares to service quotas. You can also set alarms so that you will be notified when you approach a quota.

To visualize a service quota and optionally set an alarm

1. Open the CloudWatch console at https://console.amazonaws.cn/cloudwatch/.
2. In the navigation pane, choose **Metrics**.
3. On the **All metrics** tab, choose **Usage**, then choose **By Amazon Resource**.

   The list of service quota usage metrics appears.
4. Select the check box next to one of the metrics.

   The graph displays your current usage of that Amazon resource.
5. To add your service quota to the graph, do the following:

   a. Choose the **Graphed metrics** tab.

   b. Choose **Math expression**, **Start with an empty expression**. Then in the new row, under **Details**, enter **`SERVICE_QUOTA(m1)`**.

      A new line is added to the graph, displaying the service quota for the resource represented in the metric.
6. To see your current usage as a percentage of the quota, add a new expression or change the current **SERVICE_QUOTA** expression. For the new expression, use **`m1/60/SERVICE_QUOTA(m1)*100`**.
7. (Optional) To set an alarm that notifies you if you approach the service quota, do the following:

   a. On the **`m1/60/SERVICE_QUOTA(m1)*100`** row, under **Actions**, choose the alarm icon. It looks like a bell.

The alarm creation page appears.

b. Under **Conditions**, ensure that **Threshold type** is **Static** and **Whenever Expression1 is** is set to **Greater**. Under **than**, enter **80**. This creates an alarm that goes into ALARM state when your usage exceeds 80 percent of the quota.

c. Choose **Next**.

d. On the next page, select an Amazon SNS topic or create a new one. This topic is notified when the alarm goes to ALARM state. Then choose **Next**.

e. On the next page, enter a name and description for the alarm, and then choose **Next**.

f. Choose **Create alarm**.

Amazon ECR usage metrics

You can use CloudWatch usage metrics to provide visibility into your account's usage of resources. Use these metrics to visualize your current service usage on CloudWatch graphs and dashboards.

Amazon ECR usage metrics correspond to Amazon service quotas. You can configure alarms that alert you when your usage approaches a service quota. For more information about Amazon ECR service quotas, see Amazon ECR service quotas.

Amazon ECR publishes the following metrics in the `Amazon/Usage` namespace.

Metric	Description
`CallCount`	The number of API action calls from your account. The resources are defined by the dimensions associated with the metric. The most useful statistic for this metric is `SUM`, which represent s the sum of the values from all contributors during the period defined.

The following dimensions are used to refine the usage metrics that are published by Amazon ECR.

Dimension	Description
`Service`	The name of the Amazon service containing the resource. For Amazon ECR usage metrics, the value for this dimension is `ECR`.
`Type`	The type of entity that is being reported. Currently, the only valid value for Amazon ECR usage metrics is `API`.
`Resource`	The type of resource that is running. Currently, Amazon ECR returns information on your API usage for the following API actions. • GetAuthorizationToken • BatchCheckLayerAvailability • InitiateLayerUpload • UploadLayerPart • CompleteLayerUpload • PutImage • BatchGetImage • GetDownloadUrlForLayer
`Class`	The class of resource being tracked. Currently, Amazon ECR does not use the class dimension.

Amazon ECR usage reports

Amazon provides a free reporting tool called Cost Explorer that enables you to analyze the cost and usage of your Amazon ECR resources.

Use Cost Explorer to view charts of your usage and costs. You can view data from the previous 13 months and forecast how much you are likely to spend for the next three months. You can use Cost Explorer to see patterns in how much you spend on Amazon resources over time, identify areas that need further inquiry, and see trends that you can use to understand your costs. You also can specify time ranges for the data and view time data by day or by month.

The metering data in your Cost and Usage Reports shows usage across all of your Amazon ECR repositories. For more information, see Tagging your resources for billing.

For more information about creating an Amazon Cost and Usage Report, see Amazon Cost and Usage Report in the *Amazon Billing User Guide*.

Amazon ECR repository metrics

Amazon ECR sends repository pull count metrics to Amazon CloudWatch. Amazon ECR metric data is automatically sent to CloudWatch in 1-minute periods. For more information about CloudWatch, see the Amazon CloudWatch User Guide.

Topics

- Enabling CloudWatch metrics
- Available metrics and dimensions
- Viewing Amazon ECR metrics using the CloudWatch console

Enabling CloudWatch metrics

Amazon ECR sends repository metrics automatically for all repositories. There is no need to take any manual steps.

Available metrics and dimensions

The following sections list the metrics and dimensions that Amazon ECR sends to Amazon CloudWatch.

Amazon ECR metrics

Amazon ECR provides metrics for you to monitor your repositories. You can measure the pull count.

The `AWS/ECR` namespace includes the following metrics.

`RepositoryPullCount`

The total number of pulls for the images in the repository.

Valid dimensions: `RepositoryName`.

Valid statistics: Average, Minimum, Maximum, Sum, Sample Count. The most useful statistic is Sum.

Unit: Integer.

Dimensions for Amazon ECR metrics

Amazon ECR metrics use the `AWS/ECR` namespace and provide metrics for the following dimensions.

`RepositoryName`

This dimension filters the data that you request for all container images in a specified repository.

Viewing Amazon ECR metrics using the CloudWatch console

You can view Amazon ECR repository metrics on the CloudWatch console. The CloudWatch console provides a fine-grained and customizable display of your resources. For more information, see the Amazon CloudWatch User Guide.

To view metrics in the CloudWatch console

1. Open the CloudWatch console at https://console.amazonaws.cn/cloudwatch/.
2. In the left navigation, choose **Metrics**, **All metrics**.
3. On the **Browse** tab, under **Amazon Namespaces**, choose **ECR**.
4. Choose the metrics to view. Repository metrics are scoped as **ECR > Repository Metrics**.

Amazon ECR events and EventBridge

Amazon EventBridge enables you to automate your Amazon services and to respond automatically to system events such as application availability issues or resource changes. Events from Amazon services are delivered to EventBridge in near real time. You can write simple rules to indicate which events are of interest to you and include automated actions to take when an event matches a rule. The actions that can be automatically triggered include the following:

- Adding events to log groups in CloudWatch Logs

- Invoking an Amazon Lambda function
- Invoking Amazon EC2 Run Command
- Relaying the event to Amazon Kinesis Data Streams
- Activating an Amazon Step Functions state machine
- Notifying an Amazon SNS topic or an Amazon SQS queue

For more information, see Getting Started with Amazon EventBridge in the *Amazon EventBridge User Guide*.

Sample events from Amazon ECR

The following are example events from Amazon ECR. Events are emitted on a best effort basis.

Event for a completed image push

The following event is sent when each image push is completed. For more information, see Pushing a Docker image to an Amazon ECR private repository.

```
{
    "version": "0",
    "id": "13cde686-328b-6117-af20-0e5566167482",
    "detail-type": "ECR Image Action",
    "source": "aws.ecr",
    "account": "123456789012",
    "time": "2019-11-16T01:54:34Z",
    "region": "us-west-2",
    "resources": [],
    "detail": {
        "result": "SUCCESS",
        "repository-name": "my-repository-name",
        "image-digest":
 "sha256:7f5b2640fe6fb4f46592dfd3410c4a79dac4f89e4782432e0378abcd1234",
        "action-type": "PUSH",
        "image-tag": "latest"
    }
}
```

Event for a pull through cache action

The following event is sent when a pull through cache action is attempted. For more information, see Sync an upstream registry with an Amazon ECR private registry.

```
{
    "version": "0",
    "id": "85fc3613-e913-7fc4-a80c-a3753e4aa9ae",
    "detail-type": "ECR Pull Through Cache Action",
    "source": "aws.ecr",
    "account": "123456789012",
    "time": "2023-02-29T02:36:48Z",
    "region": "us-west-2",
    "resources": [
        "arn:aws:ecr:us-west-2:123456789012:repository/docker-hub/alpine"
    ],
    "detail": {
        "rule-version": "1",
        "sync-status": "SUCCESS",
        "ecr-repository-prefix": "docker-hub",
        "repository-name": "docker-hub/alpine",
        "upstream-registry-url": "public.ecr.aws",
        "image-tag": "3.17.2",
        "image-digest":
 "sha256:4aa08ef415aecc80814cb42fa41b658480779d80c77ab15EXAMPLE",
    }
}
```

Event for a completed image scan (basic scanning)

When basic scanning is enabled for your registry, the following event is sent when each image scan is completed. The `finding-severity-counts` parameter will only return a value for a severity level if one exists. For example, if the image contains no findings at `CRITICAL` level, then no critical count is returned. For more information, see Scan images for OS vulnerabilities in Amazon ECR.

> **Note**
>
> For details about events that Amazon Inspector emits when enhanced scanning is enabled, see EventBridge events sent for enhanced scanning in Amazon ECR.

```
{
```

```
    "version": "0",
    "id": "85fc3613-e913-7fc4-a80c-a3753e4aa9ae",
    "detail-type": "ECR Image Scan",
    "source": "aws.ecr",
    "account": "123456789012",
    "time": "2019-10-29T02:36:48Z",
    "region": "us-east-1",
    "resources": [
        "arn:aws:ecr:us-east-1:123456789012:repository/my-repository-name"
    ],
    "detail": {
        "scan-status": "COMPLETE",
        "repository-name": "my-repository-name",
        "finding-severity-counts": {
        "CRITICAL": 10,
        "MEDIUM": 9
      },
        "image-digest":
 "sha256:7f5b2640fe6fb4f46592dfd3410c4a79dac4f89e4782432e0378abcd1234",
        "image-tags": []
    }
}
```

Event for a change notification on a resource with enhanced scanning enabled (enhanced scanning)

When enhanced scanning is enabled for your registry, the following event is sent by Amazon ECR when there is a change with a resource that has enhanced scanning enabled. This includes new repositories being created, the scan frequency for a repository being changed, or when images are created or deleted in repositories with enhanced scanning enabled. For more information, see Scan images for software vulnerabilities in Amazon ECR.

```
{
 "version": "0",
 "id": "0c18352a-a4d4-6853-ef53-0ab8638973bf",
 "detail-type": "ECR Scan Resource Change",
 "source": "aws.ecr",
 "account": "123456789012",
 "time": "2021-10-14T20:53:46Z",
 "region": "us-east-1",
 "resources": [],
 "detail": {
  "action-type": "SCAN_FREQUENCY_CHANGE",
```

```
  "repositories": [{
    "repository-name": "repository-1",
    "repository-arn": "arn:aws:ecr:us-east-1:123456789012:repository/repository-1",
    "scan-frequency": "SCAN_ON_PUSH",
    "previous-scan-frequency": "MANUAL"
   },
   {
    "repository-name": "repository-2",
    "repository-arn": "arn:aws:ecr:us-east-1:123456789012:repository/repository-2",
    "scan-frequency": "CONTINUOUS_SCAN",
    "previous-scan-frequency": "SCAN_ON_PUSH"
   },
   {
    "repository-name": "repository-3",
    "repository-arn": "arn:aws:ecr:us-east-1:123456789012:repository/repository-3",
    "scan-frequency": "CONTINUOUS_SCAN",
    "previous-scan-frequency": "SCAN_ON_PUSH"
   }
  ],
  "resource-type": "REPOSITORY",
  "scan-type": "ENHANCED"
 }
}
```

Event for an image deletion

The following event is sent when an image is deleted. For more information, see Deleting an image in Amazon ECR.

```
{
    "version": "0",
    "id": "dd3b46cb-2c74-f49e-393b-28286b67279d",
    "detail-type": "ECR Image Action",
    "source": "aws.ecr",
    "account": "123456789012",
    "time": "2019-11-16T02:01:05Z",
    "region": "us-west-2",
    "resources": [],
    "detail": {
        "result": "SUCCESS",
        "repository-name": "my-repository-name",
        "image-digest":
 "sha256:7f5b2640fe6fb4f46592dfd3410c4a79dac4f89e4782432e0378abcd1234",
```

```
        "action-type": "DELETE",
        "image-tag": "latest"
    }
}
```

Logging Amazon ECR actions with Amazon CloudTrail

Amazon ECR is integrated with Amazon CloudTrail, a service that provides a record of actions taken by a user, a role, or an Amazon service in Amazon ECR. CloudTrail captures the following Amazon ECR actions as events:

- All API calls, including calls from the Amazon ECR console
- All actions taken due to the encryption settings on your repositories
- All actions taken due to lifecycle policy rules, including both successful and unsuccessful actions

> **Important**
>
> Due to the size limitations of individual CloudTrail events, for lifecycle policy actions where 10 or more images are expired Amazon ECR sends multiple events to CloudTrail. Additionally, Amazon ECR includes a maximum of 100 tags per image.

When a trail is created, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon ECR. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using this information, you can determine the request that was made to Amazon ECR, the originating IP address, who made the request, when it was made, and additional details.

For more information, see the Amazon CloudTrail User Guide.

Amazon ECR information in CloudTrail

CloudTrail is enabled on your Amazon account when you create the account. When activity occurs in Amazon ECR, that activity is recorded in a CloudTrail event along with other Amazon service events in **Event history**. You can view, search, and download recent events in your Amazon account. For more information, see Viewing Events with CloudTrail Event History.

For an ongoing record of events in your Amazon account, including events for Amazon ECR, create a trail. A trail enables CloudTrail to deliver log files to an Amazon S3 bucket. When you create

a trail in the console, you can apply the trail to a single Region or to all Regions. The trail logs events in the Amazon partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other Amazon services to analyze and act upon the event data collected in CloudTrail logs. For more information, see:

- Creating a trail for your Amazon account
- Amazon service integrations with CloudTrail logs
- Configuring Amazon SNS notifications for CloudTrail
- Receiving CloudTrail log files from multiple Regions and Receiving CloudTrail log files from multiple accounts

All Amazon ECR API actions are logged by CloudTrail and are documented in the Amazon Elastic Container Registry API Reference. When you perform common tasks, sections are generated in the CloudTrail log files for each API action that is part of that task. For example, when you create a repository, `GetAuthorizationToken`, `CreateRepository` and `SetRepositoryPolicy` sections are generated in the CloudTrail log files. When you push an image to a repository, `InitiateLayerUpload`, `UploadLayerPart`, `CompleteLayerUpload`, and `PutImage` sections are generated. When you pull an image, `GetDownloadUrlForLayer` and `BatchGetImage` sections are generated. For examples of these common tasks, see CloudTrail log entry examples.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or user credentials
- Whether the request was made with temporary security credentials for a role or federated user
- Whether the request was made by another Amazon service

For more information, see the CloudTrail `userIdentity` Element.

Understanding Amazon ECR log file entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and other information. CloudTrail log files are not an ordered stack trace of the public API calls, so they do not appear in any specific order.

CloudTrail log entry examples

The following are CloudTrail log entry examples for a few common Amazon ECR tasks.

> **Note**
>
> These examples have been formatted for improved readability. In a CloudTrail log file, all entries and events are concatenated into a single line. In addition, this example has been limited to a single Amazon ECR entry. In a real CloudTrail log file, you see entries and events from multiple Amazon services.

Topics

- Example: Create repository action
- Example: Amazon KMS CreateGrant API action when creating an Amazon ECR repository
- Example: Image push action
- Example: Image pull action
- Example: Image lifecycle policy action

Example: Create repository action

The following example shows a CloudTrail log entry that demonstrates the `CreateRepository` action.

```
{
    "eventVersion": "1.04",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE:account_name",
        "arn": "arn:aws:sts::123456789012:user/Mary_Major",
        "accountId": "123456789012",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "sessionContext": {
            "attributes": {
                "mfaAuthenticated": "false",
                "creationDate": "2018-07-11T21:54:07Z"
            },
            "sessionIssuer": {
                "type": "Role",
```

```
                "principalId": "AIDACKCEVSQ6C2EXAMPLE",
                "arn": "arn:aws:iam::123456789012:role/Admin",
                "accountId": "123456789012",
                "userName": "Admin"
            }
        }
    },
    "eventTime": "2018-07-11T22:17:43Z",
    "eventSource": "ecr.amazonaws.com",
    "eventName": "CreateRepository",
    "awsRegion": "us-east-2",
    "sourceIPAddress": "203.0.113.12",
    "userAgent": "console.amazonaws.com",
    "requestParameters": {
        "repositoryName": "testrepo"
    },
    "responseElements": {
        "repository": {
            "repositoryArn": "arn:aws:ecr:us-east-2:123456789012:repository/testrepo",
            "repositoryName": "testrepo",
            "repositoryUri": "123456789012.dkr.ecr.us-east-2.amazonaws.com/testrepo",
            "createdAt": "Jul 11, 2018 10:17:44 PM",
            "registryId": "123456789012"
        }
    },
    "requestID": "cb8c167e-EXAMPLE",
    "eventID": "e3c6f4ce-EXAMPLE",
    "resources": [
        {
            "ARN": "arn:aws:ecr:us-east-2:123456789012:repository/testrepo",
            "accountId": "123456789012"
        }
    ],
    "eventType": "AwsApiCall",
    "recipientAccountId": "123456789012"
}
```


Example: Amazon KMS CreateGrant API action when creating an Amazon ECR repository

The following example shows a CloudTrail log entry that demonstrates the Amazon KMS CreateGrant action when creating an Amazon ECR repository with KMS encryption enabled. For each repository that is created with KMS encryption is enabled, you should see two CreateGrant log entries in CloudTrail.

```
{
    "eventVersion": "1.05",
    "userIdentity": {
        "type": "IAMUser",
        "principalId": "AIDAIEP6W46J43IG7LXAQ",
        "arn": "arn:aws:iam::123456789012:user/Mary_Major",
        "accountId": "123456789012",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "userName": "Mary_Major",
        "sessionContext": {
            "sessionIssuer": {

            },
            "webIdFederationData": {

            },
            "attributes": {
                "mfaAuthenticated": "false",
                "creationDate": "2020-06-10T19:22:10Z"
            }
        },
        "invokedBy": "AWS Internal"
    },
    "eventTime": "2020-06-10T19:22:10Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "CreateGrant",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "203.0.113.12",
    "userAgent": "console.amazonaws.com",
    "requestParameters": {
        "keyId": "4b55e5bf-39c8-41ad-b589-18464af7758a",
        "granteePrincipal": "ecr.us-west-2.amazonaws.com",
        "operations": [
            "GenerateDataKey",
            "Decrypt"
        ],
        "retiringPrincipal": "ecr.us-west-2.amazonaws.com",
        "constraints": {
            "encryptionContextSubset": {
                "aws:ecr:arn": "arn:aws:ecr:us-west-2:123456789012:repository/testrepo"
            }
        }
    },
```

```
    "responseElements": {
        "grantId": "3636af9adfee1accb67b83941087dcd45e7fadc4e74ff0103bb338422b5055f3"
    },
    "requestID": "047b7dea-b56b-4013-87e9-a089f0f6602b",
    "eventID": "af4c9573-c56a-4886-baca-a77526544469",
    "readOnly": false,
    "resources": [
        {
            "accountId": "123456789012",
            "type": "AWS::KMS::Key",
            "ARN": "arn:aws:kms:us-west-2:123456789012:key/4b55e5bf-39c8-41ad-
b589-18464af7758a"
        }
    ],
    "eventType": "AwsApiCall",
    "recipientAccountId": "123456789012"
}
```

Example: Image push action

The following example shows a CloudTrail log entry that demonstrates an image push which uses the `PutImage` action.

> **Note**
>
> When pushing an image, you will also see `InitiateLayerUpload`, `UploadLayerPart`, and `CompleteLayerUpload` references in the CloudTrail logs.

```
{
    "eventVersion": "1.04",
    "userIdentity": {
    "type": "IAMUser",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE:account_name",
    "arn": "arn:aws:sts::123456789012:user/Mary_Major",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
  "userName": "Mary_Major",
  "sessionContext": {
   "attributes": {
    "mfaAuthenticated": "false",
    "creationDate": "2019-04-15T16:42:14Z"
```

```
    }
   }
  },
  "eventTime": "2019-04-15T16:45:00Z",
  "eventSource": "ecr.amazonaws.com",
  "eventName": "PutImage",
  "awsRegion": "us-east-2",
  "sourceIPAddress": "203.0.113.12",
  "userAgent": "console.amazonaws.com",
  "requestParameters": {
   "repositoryName": "testrepo",
   "imageTag": "latest",
   "registryId": "123456789012",
   "imageManifest": "{\n   \"schemaVersion\": 2,\n   \"mediaType\": \"application/
vnd.docker.distribution.manifest.v2+json\",\n   \"config\": {\n      \"mediaType\":
 \"application/vnd.docker.container.image.v1+json\",\n      \"size\": 5543,\n
 \"digest\": \"sha256:000b9b805af1cdb60628898c9f411996301a1c13afd3dbef1d8a16ac6dbf503a
\"\n   },\n   \"layers\": [\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 43252507,\n
 \"digest\": \"sha256:3b37166ec61459e76e33282dda08f2a9cd698ca7e3d6bc44e6a6e7580cdeff8e
\"\n      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 846,\n         \"digest
\": \"sha256:504facff238fde83f1ca8f9f54520b4219c5b8f80be9616ddc52d31448a044bd
\"\n      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 615,\n         \"digest
\": \"sha256:ebbcacd28e101968415b0c812b2d2dc60f969e36b0b08c073bf796e12b1bb449\"\n
      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 850,\n         \"digest
\": \"sha256:c7fb3351ecad291a88b92b600037e2435c84a347683d540042086fe72c902b8a
\"\n      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 168,\n         \"digest\":
 \"sha256:2e3debadcbf7e542e2aefbce1b64a358b1931fb403b3e4aeca27cb4d809d56c2\"\n      },
\n      {\n         \"mediaType\": \"application/vnd.docker.image.rootfs.diff.tar.gzip
\",\n         \"size\": 37720774,\n         \"digest\":
 \"sha256:f8c9f51ad524d8ae9bf4db69cd3e720ba92373ec265f5c390ffb21bb0c277941\"\n
      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 30432107,\n
 \"digest\": \"sha256:813a50b13f61cf1f8d25f19fa96ad3aa5b552896c83e86ce413b48b091d7f01b
\"\n      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 197,\n         \"digest
\": \"sha256:7ab043301a6187ea3293d80b30ba06c7bf1a0c3cd4c43d10353b31bc0cecfe7d
\"\n      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 154,\n         \"digest
\": \"sha256:67012cca8f31dc3b8ee2305e7762fee20c250513effdedb38a1c37784a5a2e71\"\n
```

```
      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 176,\n         \"digest
\": \"sha256:3bc892145603fffc9b1c97c94e2985b4cb19ca508750b15845a5d97becbd1a0e
\"\n      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 183,\n         \"digest
\": \"sha256:6f1c79518f18251d35977e7e46bfa6c6b9cf50df2a79d4194941d95c54258d18\"\n
      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 212,\n         \"digest
\": \"sha256:b7bcfbc2e2888afebede4dd1cd5eebf029bb6315feeaf0b56e425e11a50afe42\"\n
      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 212,\n         \"digest\":
 \"sha256:2b220f8b0f32b7c2ed8eaafe1c802633bbd94849b9ab73926f0ba46cdae91629\"\n      }\n
   ]\n}"
 },
 "responseElements": {
  "image": {
   "repositoryName": "testrepo",
   "imageManifest": "{\n   \"schemaVersion\": 2,\n   \"mediaType\": \"application/
vnd.docker.distribution.manifest.v2+json\",\n   \"config\": {\n      \"mediaType\":
 \"application/vnd.docker.container.image.v1+json\",\n      \"size\": 5543,\n
 \"digest\": \"sha256:000b9b805af1cdb60628898c9f411996301a1c13afd3dbef1d8a16ac6dbf503a
\"\n   },\n   \"layers\": [\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 43252507,\n
 \"digest\": \"sha256:3b37166ec61459e76e33282dda08f2a9cd698ca7e3d6bc44e6a6e7580cdeff8e
\"\n      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 846,\n         \"digest
\": \"sha256:504facff238fde83f1ca8f9f54520b4219c5b8f80be9616ddc52d31448a044bd
\"\n      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 615,\n         \"digest
\": \"sha256:ebbcacd28e101968415b0c812b2d2dc60f969e36b0b08c073bf796e12b1bb449\"\n
      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 850,\n         \"digest
\": \"sha256:c7fb3351ecad291a88b92b600037e2435c84a347683d540042086fe72c902b8a
\"\n      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 168,\n         \"digest\":
 \"sha256:2e3debadcbf7e542e2aefbce1b64a358b1931fb403b3e4aeca27cb4d809d56c2\"\n      },
\n      {\n         \"mediaType\": \"application/vnd.docker.image.rootfs.diff.tar.gzip
\",\n         \"size\": 37720774,\n         \"digest\":
 \"sha256:f8c9f51ad524d8ae9bf4db69cd3e720ba92373ec265f5c390ffb21bb0c277941\"\n
      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 30432107,\n
 \"digest\": \"sha256:813a50b13f61cf1f8d25f19fa96ad3aa5b552896c83e86ce413b48b091d7f01b
\"\n      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 197,\n         \"digest
```

```
\": \"sha256:7ab043301a6187ea3293d80b30ba06c7bf1a0c3cd4c43d10353b31bc0cecfe7d
\"\n      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 154,\n         \"digest
\": \"sha256:67012cca8f31dc3b8ee2305e7762fee20c250513effdedb38a1c37784a5a2e71\"\n
      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 176,\n         \"digest
\": \"sha256:3bc892145603fffc9b1c97c94e2985b4cb19ca508750b15845a5d97becbd1a0e
\"\n      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 183,\n         \"digest
\": \"sha256:6f1c79518f18251d35977e7e46bfa6c6b9cf50df2a79d4194941d95c54258d18\"\n
      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 212,\n         \"digest
\": \"sha256:b7bcfbc2e2888afebede4dd1cd5eebf029bb6315feeaf0b56e425e11a50afe42\"\n
      },\n      {\n         \"mediaType\": \"application/
vnd.docker.image.rootfs.diff.tar.gzip\",\n         \"size\": 212,\n         \"digest\":
 \"sha256:2b220f8b0f32b7c2ed8eaafe1c802633bbd94849b9ab73926f0ba46cdae91629\"\n      }\n
   ]\n}",
   "registryId": "123456789012",
   "imageId": {
    "imageDigest":
 "sha256:98c8b060c21d9adbb6b8c41b916e95e6307102786973ab93a41e8b86d1fc6d3e",
    "imageTag": "latest"
   }
  }
 },
 "requestID": "cf044b7d-5f9d-11e9-9b2a-95983139cc57",
 "eventID": "2bfd4ee2-2178-4a82-a27d-b12939923f0f",
 "resources": [{
  "ARN": "arn:aws:ecr:us-east-2:123456789012:repository/testrepo",
  "accountId": "123456789012"
 }],
 "eventType": "AwsApiCall",
 "recipientAccountId": "123456789012"
}
```


Example: Image pull action

The following example shows a CloudTrail log entry that demonstrates an image pull which uses the `BatchGetImage` action.

> **Note**
>
> When pulling an image, if you don't already have the image locally, you will also see `GetDownloadUrlForLayer` references in the CloudTrail logs.


```
{
    "eventVersion": "1.04",
    "userIdentity": {
    "type": "IAMUser",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE:account_name",
    "arn": "arn:aws:sts::123456789012:user/Mary_Major",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
  "userName": "Mary_Major",
  "sessionContext": {
   "attributes": {
    "mfaAuthenticated": "false",
    "creationDate": "2019-04-15T16:42:14Z"
   }
  }
 },
 "eventTime": "2019-04-15T17:23:20Z",
 "eventSource": "ecr.amazonaws.com",
 "eventName": "BatchGetImage",
 "awsRegion": "us-east-2",
 "sourceIPAddress": "203.0.113.12",
 "userAgent": "console.amazonaws.com",
 "requestParameters": {
  "imageIds": [{
   "imageTag": "latest"
  }],
  "acceptedMediaTypes": [
   "application/json",
   "application/vnd.oci.image.manifest.v1+json",
   "application/vnd.oci.image.index.v1+json",
   "application/vnd.docker.distribution.manifest.v2+json",
   "application/vnd.docker.distribution.manifest.list.v2+json",
   "application/vnd.docker.distribution.manifest.v1+prettyjws"
  ],
  "repositoryName": "testrepo",
  "registryId": "123456789012"
```

```
 },
 "responseElements": null,
 "requestID": "2a1b97ee-5fa3-11e9-a8cd-cd2391aeda93",
 "eventID": "c84f5880-c2f9-4585-9757-28fa5c1065df",
 "resources": [{
  "ARN": "arn:aws:ecr:us-east-2:123456789012:repository/testrepo",
  "accountId": "123456789012"
 }],
 "eventType": "AwsApiCall",
 "recipientAccountId": "123456789012"
}
```

Example: Image lifecycle policy action

The following example shows a CloudTrail log entry that demonstrates when an image is expired due to a lifecycle policy rule. This event type can be located by filtering for `PolicyExecutionEvent` for the event name field.

> **Important**
>
> Due to the size limitations of individual CloudTrail events, for lifecycle policy actions where 10 or more images are expired Amazon ECR sends multiple events to CloudTrail. Additionally, Amazon ECR includes a maximum of 100 tags per image.

```
{
    "eventVersion": "1.05",
    "userIdentity": {
        "accountId": "123456789012",
        "invokedBy": "AWS Internal"
    },
    "eventTime": "2020-03-12T20:22:12Z",
    "eventSource": "ecr.amazonaws.com",
    "eventName": "PolicyExecutionEvent",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "AWS Internal",
    "userAgent": "AWS Internal",
    "requestParameters": null,
    "responseElements": null,
    "eventID": "9354dd7f-9aac-4e9d-956d-12561a4923aa",
    "readOnly": true,
    "resources": [
```

```
        {
            "ARN": "arn:aws:ecr:us-west-2:123456789012:repository/testrepo",
            "accountId": "123456789012",
            "type": "AWS::ECR::Repository"
        }
    ],
    "eventType": "AwsServiceEvent",
    "recipientAccountId": "123456789012",
    "serviceEventDetails": {
        "repositoryName": "testrepo",
        "lifecycleEventPolicy": {
            "lifecycleEventRules": [
                {
                    "rulePriority": 1,
                    "description": "remove all images > 2",
                    "lifecycleEventSelection": {
                        "tagStatus": "Any",
                        "tagPrefixList": [],
                        "countType": "Image count more than",
                        "countNumber": 2
                    },
                    "action": "expire"
                }
            ],
            "lastEvaluatedAt": 0,
            "policyVersion": 1,
            "policyId": "ceb86829-58e7-9498-920c-aa042e33037b"
        },
        "lifecycleEventImageActions": [
            {
                "lifecycleEventImage": {
                    "digest":
 "sha256:ddba4d27a7ffc3f86dd6c2f92041af252a1f23a8e742c90e6e1297bfa1bc0c45",
                    "tagStatus": "Tagged",
                    "tagList": [
                        "alpine"
                    ],
                    "pushedAt": 1584042813000
                },
                "rulePriority": 1
            },
            {
                "lifecycleEventImage": {
```

```
                    "digest":
 "sha256:6ab380c5a5acf71c1b6660d645d2cd79cc8ce91b38e0352cbf9561e050427baf",
                    "tagStatus": "Tagged",
                    "tagList": [
                        "centos"
                    ],
                    "pushedAt": 1584042842000
                },
                "rulePriority": 1
            }
        ]
    }
}
```

Amazon ECR service quotas

The following table provides the default service quotas for Amazon Elastic Container Registry (Amazon ECR).

Name	Default	Adjusted	Description
Filters per rule in a replication configuration	Each supported Region: 100	No	The maximum number of filters per rule in a replication configuration.
Images per repository	Each supported Region: 10,000	Yes	The maximum number of images per repository.
Layer parts	Each supported Region: 4,200	No	The maximum number of layer parts. This is only applicable if you are using Amazon ECR API actions directly to initiate multipart uploads for image push operations.
Lifecycle policy length	Each supported Region: 30,720	No	The maximum number of characters in a lifecycle policy.
Maximum layer part size	Each supported Region: 10	No	The maximum size (MiB) of a layer part. This is only applicable if you are using Amazon ECR API actions directly to initiate multipart uploads for image push operations.
Maximum layer size	Each supported Region: 52,000	No	The maximum size (MiB) of a layer.

Name	Default	Adjuse	Description
Minimum layer part size	Each supported Region: 5	No	The minimum size (MiB) of a layer part. This is only applicable if you are using Amazon ECR API actions directly to initiate multipart uploads for image push operations.
Pull through cache rules per registry	Each supported Region: 50	No	The maximum number of pull-through cache rules.
Rate of BatchCheckLayerAvailability requests	Each supported Region: 1,000 per second	Yes	The maximum number of BatchCheckLayerAvailability requests that you can make per second in the current Region. When an image is pushed to a repository, each image layer is checked to verify if it has been uploaded before. If it has been uploaded, then the image layer is skipped.

Name	Default	Adjusue	Description
Rate of BatchGetImage requests	Each supported Region: 2,000 per second	Yes	The maximum number of BatchGetImage requests that you can make per second in the current Region. When an image is pulled, the BatchGetImage API is called once to retrieve the image manifest. If you request a quota increase for this API, review your GetDownloadUrlForLayer usage as well.
Rate of CompleteLayerUpload requests	Each supported Region: 100 per second	Yes	The maximum number of CompleteLayerUpload requests that you can make per second in the current Region. When an image is pushed, the CompleteLayerUpload API is called once per each new image layer to verify that the upload has completed.
Rate of GetAuthorizationToken requests	Each supported Region: 500 per second	Yes	The maximum number of GetAuthorizationToken requests that you can make per second in the current Region.

Name	Default	Adjusue	Description
Rate of GetDownloadUrlForLayer requests	Each supported Region: 3,000 per second	Yes	The maximum number of GetDownloadUrlForLayer requests that you can make per second in the current Region. When an image is pulled, the GetDownloadUrlForLayer API is called once per image layer that is not already cached. If you request a quota increase for this API, review your BatchGetImage usage as well.
Rate of InitiateLayerUpload requests	Each supported Region: 100 per second	Yes	The maximum number of InitiateLayerUpload requests that you can make per second in the current Region. When an image is pushed, the InitiateLayerUploa d API is called once per image layer that has not already been uploaded. Whether or not an image layer has been uploaded is determined by the BatchCheckLayerAva ilability API action.

Name	Default	Adjuse	Description
Rate of PutImage requests	Each supported Region: 10 per second	Yes	The maximum number of PutImage requests that you can make per second in the current Region. When an image is pushed and all new image layers have been uploaded, the PutImage API is called once to create or update the image manifest and the tags associated with the image.
Rate of UploadLayerPart requests	Each supported Region: 500 per second	Yes	The maximum number of UploadLayerPart requests that you can make per second in the current Region. When an image is pushed, each new image layer is uploaded in parts and the UploadLayerPart API is called once per each new image layer part.
Rate of image scans	Each supported Region: 1	No	The maximum number of image scans per image, per 24 hours.
Registered repositories	Each supported Region: 10,000	Yes	The maximum number of repositories that you can create in this account in the current Region.

Name	Default	Adjus e	Description
Rules per lifecycle policy	Each supported Region: 50	No	The maximum number of rules in a lifecycle policy
Rules per replication configuration	Each supported Region: 10	No	The maximum number of rules in a replication configuration.
Tags per image	Each supported Region: 1,000	No	The maximum number of tags per image.
Unique destinations across all rules in a replication configuration	Each supported Region: 25	No	The maximum number of unique destinati ons across all rules in a replication configuration.

Amazon ECR troubleshooting

This chapter helps you find diagnostic information for Amazon ECR, and provides troubleshooting steps for common issues and error messages.

Topics

- Troubleshooting Docker commands and issues when using Amazon ECR
- Troubleshooting Amazon ECR error messages

Troubleshooting Docker commands and issues when using Amazon ECR

In some cases, running a Docker command against Amazon ECR might result in an error message. Some common error messages and potential solutions are explained below.

Topics

- Docker logs do not contain expected error messages
- Error: "Filesystem Verification Failed" or "404: Image Not Found" when pulling an image from an Amazon ECR repository
- Error: "Filesystem Layer Verification Failed" when pulling images from Amazon ECR
- HTTP 403 Errors or "no basic auth credentials" error when pushing to repository

Docker logs do not contain expected error messages

To begin debugging any Docker-related issue, start by turning on Docker debugging output on the Docker daemon running on your host instances. If you are using images pulled from Amazon ECR on Amazon ECS container instances, see Configuring verbose output from the Docker daemon in the *Amazon Elastic Container Service Developer Guide*.

Error: "Filesystem Verification Failed" or "404: Image Not Found" when pulling an image from an Amazon ECR repository

You may receive the error `Filesystem verification failed` when using the **docker pull** command to pull an image from an Amazon ECR repository with Docker 1.9 or above. You may receive the error `404: Image not found` when you are using Docker versions before 1.9.

Some possible reasons and their explanations are given below.

The local disk is full

If the local disk on which you're running **docker pull** is full, then the SHA-1 hash calculated on the local file may be different than the one calculated by Amazon ECR. Check that your local disk has enough remaining free space to store the Docker image you are pulling. You can also delete old images to make room for new ones. Use the **docker images** command to see a list of all locally downloaded Docker images, along with their sizes.

Client cannot connect to the remote repository due to network error

Calls to an Amazon ECR repository require a functioning connection to the internet. Verify your network settings, and verify that other tools and applications can access resources on the internet. If you are running **docker pull** on an Amazon EC2 instance in a private subnet, verify that the subnet has a route to the internet. Use a network address translation (NAT) server or a managed NAT gateway.

Currently, calls to an Amazon ECR repository also require network access through your corporate firewall to Amazon Simple Storage Service (Amazon S3). If your organization uses firewall software or a NAT device that allows service endpoints, ensure that the Amazon S3 service endpoints for your current Region are allowed.

If you are using Docker behind an HTTP proxy, you can configure Docker with the appropriate proxy settings. For more information, see HTTP proxy in the Docker documentation.

Error: "Filesystem Layer Verification Failed" when pulling images from Amazon ECR

You may receive the error `image image-name not found` when pulling images using the **docker pull** command. If you inspect the Docker logs, you may see an error like the following:

```
filesystem layer verification failed for digest sha256:2b96f...
```

This error indicates that one or more of the layers for your image has failed to download. Some possible reasons and their explanations are given below.

You are using an older version of Docker

This error can occur in a small percentage of cases when using a Docker version less than 1.10. Upgrade your Docker client to 1.10 or greater.

Your client has encountered a network or disk error

A full disk or a network issue may prevent one or more layers from downloading, as discussed earlier about the `Filesystem verification failed` message. Follow the recommendations above to ensure that your filesystem is not full, and that you have enabled access to Amazon S3 from within your network.

HTTP 403 Errors or "no basic auth credentials" error when pushing to repository

There are times when you may receive an `HTTP 403 (Forbidden)` error, or the error message `no basic auth credentials` from the **docker push** or **docker pull** commands, even if you have successfully authenticated to Docker using the **aws ecr get-login-password** command. The following are some known causes of this issue:

You have authenticated to a different region

Authentication requests are tied to specific regions, and cannot be used across regions. For example, if you obtain an authorization token from US West (Oregon), you cannot use it to authenticate against your repositories in US East (N. Virginia). To resolve the issue, ensure that you have retrieved an authentication token from the same Region your repository exists in. For more information, see the section called “Registry authentication”.

You have authenticated to push to a repository you don't have permissions for

You do not have the necessary permissions to push to the repository. For more information, see Private repository policies in Amazon ECR.

Your token has expired

The default authorization token expiration period for tokens obtained using the `GetAuthorizationToken` operation is 12 hours.

Bug in `wincred` credential manager

Some versions of Docker for Windows use a credential manager called `wincred`, which does not properly handle the Docker login command produced by **aws ecr get-login-password** (for more information, see https://github.com/docker/docker/issues/22910). You can run the Docker login command that is output, but when you try to push or pull images, those commands fail. You can work around this bug by removing the `https://` scheme from the registry argument in the Docker login command that is output from **aws ecr get-login-password**. An example Docker login command without the HTTPS scheme is shown below.

```
docker login -u AWS -p <password> <aws_account_id>.dkr.ecr.<region>.amazonaws.com
```

Troubleshooting Amazon ECR error messages

In some cases, an API call that you have initiated through the Amazon ECR console or the Amazon CLI exits with an error message. Some common error messages and potential solutions are explained below.

HTTP 429: Too Many Requests or ThrottleException

You may receive a `429: Too Many Requests` error or a `ThrottleException` error from one or more Amazon ECR actions or API calls. This indicates that you are calling a single endpoint in Amazon ECR repeatedly over a short interval, and that your requests are getting throttled. Throttling occurs when calls to a single endpoint from a single user exceed a certain threshold over a period of time.

Each API operations in Amazon ECR has a rate throttles associated with it. For example, the throttle for the `GetAuthorizationToken` action is 20 transaction per second (TPS), with up to a 200 TPS burst allowed. In each region, each account receives a bucket that can store up to 200 `GetAuthorizationToken` credits. These credits are replenished at a rate of 20 per second. If your bucket has 200 credits, you could achieve 200 `GetAuthorizationToken` API transactions per second for one second, and then sustain 20 transactions per second indefinitely. For more information on the rate limits for Amazon ECR APIs, see Amazon ECR service quotas.

To handle throttling errors, implement a retry function with incremental backoff into your code. For more information, see Retry behavior in the *Amazon SDKs and Tools Reference Guide*. Another option is to request a rate limit increase, which you can do using the Service Quotas console. .

HTTP 403: "User [arn] is not authorized to perform [operation]"

You may receive the following error when attempting to perform an action with Amazon ECR:

```
$ aws ecr get-login-password
A client error (AccessDeniedException) occurred when calling the GetAuthorizationToken
 operation:
    User: arn:aws:iam::account-number:user/username is not authorized to perform:
    ecr:GetAuthorizationToken on resource: *
```

This indicates that your user does not have permissions granted to use Amazon ECR, or that those permissions are not set up correctly. In particular, if you are performing actions against an Amazon ECR repository, verify that the user has been granted permissions to access that repository. For more information about creating and verifying permissions for Amazon ECR, see Identity and Access Management for Amazon Elastic Container Registry.

HTTP 404: "Repository Does Not Exist" error

If you specify a Docker Hub repository that does not currently exist, Docker Hub creates it automatically. With Amazon ECR, new repositories must be explicitly created before they can be used. This prevents new repositories from being created accidentally (for example, due to typos), and it also ensures that an appropriate security access policy is explicitly assigned to any new repositories. For more information about creating repositories, see Amazon ECR private repositories.

Error: Cannot perform an interactive login from a non TTY device

If you receive the error `Cannot perform an interactive login from a non TTY device`, the following troubleshooting steps should help.

- Verify that you're using Amazon CLI version 2 and that you don't have a conflicting version of Amazon CLI version 1 on your system. For more information, see Installing or updating the latest version of the Amazon CLI.
- Verify that you've configured your Amazon CLI with valid credentials. For more information, see Installing or updating the latest version of the Amazon CLI.

- Verify that the syntax of your Amazon CLI command is correct.

Document history

The following table describes the important changes to the documentation since the last release of Amazon ECR. We also update the documentation frequently to address the feedback that you send us.

Change	Description	Date
Added cross-Region and cross-account replication to China Regions	Amazon ECR added support to China Region for filtering which repositories are replicated.	15 May 2024
Added GitLab container registry to pull through cache rules	Amazon ECR added support for creating pull through cache rules for the GitLab container registry.	8 May 2024
Amazon ECR lifecycle policy update to add support for using wildcards	Amazon ECR added support for wildcards in a lifecycle policy through the use of the `tagPatternList` parameter in a lifecycle policy rule. For more information, see Automate the cleanup of images by using lifecycle policies in Amazon ECR.	18 December 2023
Amazon ECR repository creation templates	Amazon ECR added support for repository creation templates. For more information, see Templates to control repositories created during a pull through cache action.	15 November 2023
Amazon ECR pull through cache added supported for authenticated upstream registries	Amazon ECR added support for using upstream registries that require authentication for your pull through cache rules. For more information, see Sync an upstream registry with an Amazon ECR private registry.	15 November 2023
AWSECRPullThroughCache_ServiceRolePolicy – Update to an existing policy	Amazon ECR added new permissions to the `AWSECRPullThroughCache_ServiceRolePolicy` policy. These permissions allow Amazon ECR to retrieve the encrypted contents of a Secrets	November 15, 2023

Change	Description	Date
	Manager secret. This is required when using a pull through cache rule to cache images from an upstream registry that requires authentication.	
Amazon ECR image signing	Amazon ECR and Amazon Signer added support for creating and pushing container image signatures using the Notary client. For more information, see Signing an image stored in an Amazon ECR private repository.	6 June 2023
Added Kubernetes container registry to pull through cache rules	Amazon ECR added support for creating pull through cache rules for the Kubernetes container registry. For more information, see Sync an upstream registry with an Amazon ECR private registry.	1 June 2023
Amazon ECR enhanced scanning duration support	Amazon Inspector added support for setting the duration that your repositories are monitored for when enhanced scanning is enabled. For more information, see Changing the enhanced scanning duration for images in Amazon Inspector.	28 June 2022
Amazon ECR sends repository pull count metrics to Amazon CloudWatch	Amazon ECR sends repository pull count metrics to Amazon CloudWatch. For more information, see Amazon ECR repository metrics.	6 January 2022
Expanded replication support	Amazon ECR added support for filtering which repositories are replicated. For more information, see Private image replication in Amazon ECR.	21 September 2021
Amazon managed policies for Amazon ECR	Amazon ECR added documentation of Amazon managed policies. For more information, see Amazon managed policies for Amazon Elastic Container Registry.	24 June 2021

Change	Description	Date
Cross-Region and cross-account replication	Amazon ECR added support for configuring replication settings for your private registry. For more information, see Private registry settings in Amazon ECR.	8 December 2020
OCI artifact support	Amazon ECR added support for pushing and pulling Open Container Initiative (OCI) artifacts. A new parameter `artifactMediaType` was added to the `DescribeImages` API response to indicate the type of artifact. For more information, see Pushing a Helm chart to an Amazon ECR private repository.	24 August 2020
Encryption at rest	Amazon ECR added support for configuring encryption for your repositories using server-side encryption with customer managed keys stored in Amazon Key Management Service (Amazon KMS). For more information, see Encryption at rest.	29 July 2020
Multi-architecture images	Amazon ECR added support for creating and pushing Docker manifest lists which are used for multi-architecture images. For more information, see Pushing a multi-architecture image to an Amazon ECR private repository.	28 April 2020
Amazon ECR Usage Metrics	Amazon ECR added CloudWatch usage metrics which provides visiblity into your account's resource usage. You also have the ability to create CloudWatch alarms from both the CloudWatch and Service Quotas consoles to get alerts when your usage approaches your applied service quota. For more information, see Amazon ECR usage metrics.	28 Feb 2020

Change	Description	Date
Updated Amazon ECR service quotas	Updated the Amazon ECR service quotas to include per-API quotas. For more information, see Amazon ECR service quotas.	19 Feb 2020
Added `get-login-password` command	Added support for **get-login-password**, which provides a simple and secure method for retrieving an authorization token. For more information, see Using an authorization token.	4 Feb 2020
Image Scanning	Added support for image scanning, which helps in identifying software vulnerabilities in your container images. Amazon ECR uses the Common Vulnerabilities and Exposures (CVEs) database from the open source CoreOS Clair project and provides you with a list of scan findings. For more information, see Scan images for software vulnerabilities in Amazon ECR.	24 Oct 2019
VPC Endpoint Policy	Added support for setting an IAM policy on the Amazon ECR interface VPC endpoints. For more information, see Create an endpoint policy for your Amazon ECR VPC endpoints.	26 Sept 2019
Image Tag Mutability	Added support for configuring a repository to be immutable to prevent image tags from being overwritten. For more information, see Preventing image tags from being overwritten in Amazon ECR.	25 July 2019

Change	Description	Date
Interface VPC Endpoints (Amazon PrivateLink)	Added support for configuring interface VPC endpoints powered by Amazon PrivateLink. This allows you to create a private connection between your VPC and Amazon ECR without requiring access over the internet, through a NAT instance, a VPN connection, or Amazon Direct Connect. For more information, see Amazon ECR interface VPC endpoints (Amazon PrivateLink).	25 Jan 2019
Resource tagging	Amazon ECR added support for adding metadata tags to your repositories. For more information, see Tagging a private repository in Amazon ECR.	18 Dec 2018
Amazon ECR Name Change	Amazon Elastic Container Registry is renamed (previously Amazon EC2 Container Registry).	21 Nov 2017
Lifecycle Policies	Amazon ECR lifecycle policies enable you to specify the lifecycle management of images in a repository. For more information, see Automate the cleanup of images by using lifecycle policies in Amazon ECR.	11 Oct 2017
Amazon ECR support for Docker image manifest 2, schema 2	Amazon ECR now supports Docker Image Manifest V2 Schema 2 (used with Docker version 1.10 and newer). For more information, see Container image manifest format support in Amazon ECR.	27 Jan 2017
Amazon ECR General Availability	Amazon Elastic Container Registry (Amazon ECR) is a managed Amazon Docker registry service that is secure, scalable, and reliable.	21 Dec 2015