aws

Developer Guide

Deadline Cloud

Deadline Cloud: Developer Guide

Table of Contents

What is AWS Deadline Cloud?

AWS Deadline Cloud is a fully-managed AWS service that enables you to have a scalable processing farm up and running in minutes. It provides an administration console for managing users, farms, queues for scheduling jobs, and fleets of workers that do the processing.

This developer guide is for pipeline, tools, and applications developers in a wide range of use cases, including the following:

- Pipeline developers and technical directors can integrate Deadline Cloud APIs and features into their custom production pipelines.
- Independent software vendors can integrate Deadline Cloud into their applications enabling digital content creation artists and users to submit Deadline Cloud render jobs seamlessly from their workstations.
- Web and cloud-based service developers can integrate Deadline Cloud rendering into their platforms, enabling customers to provide assets to view products virtually.

We provide tools that enable you to work directly with any step of your pipeline:

- A command-line interface that you can use directly or from scripts.
- The AWS SDK for 11 popular programming languages.
- A REST-based web interface that you can call from your applications.

You can also use other AWS services in your custom applications. For example, you can use:

- **AWS CloudFormation** to automate creating and removing farms, queues, and fleets.
- **Amazon CloudWatch** to gather metrics for jobs.
- **Amazon Simple Storage Service** to store and manage digital assets and job output.
- **AWS IAM Identity Center** to manage users and groups for your farms.

Open Job Description

Deadline Cloud uses the [Open Job Description (OpenJD) specification](#) to specify the details of a job. OpenJD was developed to define jobs that are portable between solutions. You use it to define a job that is a set of commands that run on worker hosts.

You can create an OpenJD job template using a submitter that Deadline Cloud provides, or you can use any tool that you want to create the template. After creating the template, you send it to Deadline Cloud. If you use a submitter, it takes care of sending the template. If you created the template another way, you call a Deadline Cloud command-line action, or you can use one of the AWS SDKs to send the job. Either way, Deadline Cloud adds the job to the specified queue and schedules the work.

Concepts and terminology for Deadline Cloud

To help you get started with AWS Deadline Cloud, this topic explains some of its key concepts and terminology.

Budget manager

Budget manager is part of the Deadline Cloud monitor. Use the budget manager to create and manage budgets. You can also use it to limit activities to stay within budget.

Deadline Cloud Client Library

The Client Library includes a command line interface and library for managing Deadline Cloud. Functionality includes submitting job bundles based on the Open Job Description specification to Deadline Cloud, downloading job attachment outputs, and monitoring your farm using the command line interface.

Digital content creation application (DCC)

Digital content creation applications (DCCs) are third-party products where you create digital content. Examples of DCCs are Maya, Nuke, and Houdini. Deadline Cloud provides job submitter integrated plugins for specific DCCs.

Farm

A farm is a where your project resources are located. It consists of queues and fleets.

Fleet

A fleet is a group of worker nodes that do the rendering. Worker nodes process jobs. A fleet can be associated to multiple queues, and a queue can be associated to multiple fleets.

Job

A job is a rendering request. Users submit jobs. Jobs contain specific job properties that are outlined as steps and tasks.

Job attachments

A job attachment is a Deadline Cloud feature that you can use to manage inputs and outputs for jobs. Job files are uploaded as job attachments during the rendering process. These files can be textures, 3D models, lighting rigs, and other similar items.

Job properties

Job properties are settings that you define when submitting a render job. Some examples include frame range, output path, job attachments, renderable camera, and more. The properties vary based on the DCC that the render is submitted from.

Job template

A job template defines the runtime environment and all processes that run as part of a Deadline Cloud job.

Queue

A queue is where submitted jobs are located and scheduled to be rendered. A queue must be associated with a fleet to create a successful render. A queue can be associated with multiple fleets.

Queue-fleet association

When a queue is associated with a fleet, there is a queue-fleet association. Use an association to schedule workers from a fleet to jobs in that queue. You can start and stop associations to control scheduling of work.

Step

A step is one particular process to run in the job.

Deadline Cloud submitter

A Deadline Cloud submitter is a digital content creation (DCC) plugin. Artists use it to submit jobs from a third-party DCC interface that they are familiar with.

Tags

A tag is a label that you can assign to an AWS resource. Each tag consists of a key and an optional value that you define.

With tags, you can categorize your AWS resources in different ways. For example, you could define a set of tags for your account's Amazon EC2 instances that help you track each instance's owner and stack level.

You can also categorize your AWS resources by purpose, owner, or environment. This approach is useful when you have many resources of the same type. You can quickly identify a specific resources based on the tags that you've assigned to it.

Task

A task is a single component of a render step.

Usage-based licensing (UBL)

Usage-based licensing (UBL) is an on-demand licensing model that is available for select third-party products. This model is pay as your go, and you are charged for the number of hours and minutes that you use.

Usage explorer

Usage explorer is a feature of Deadline Cloud monitor. It provides an approximate estimate of your costs and usage.

Worker

Workers belong to fleets and run Deadline Cloud assigned tasks to complete steps and jobs. Workers store the logs from task operations in Amazon CloudWatch Logs. Workers can also use the job attachments feature to sync inputs and outputs to an Amazon Simple Storage Service (Amazon S3) bucket.

What is a Deadline Cloud workload

With AWS Deadline Cloud, you can submit jobs to run your applications in the cloud and process data for the production of content or insights crucial to your business. Deadline Cloud uses [Open Job Description](#) (OpenJD) as the syntax for job templates, a specification designed for the needs of visual compute pipelines but applicable to many other use cases. Some example workloads include computer graphics rendering, physics simulation, and photogrammetry.

Workloads scale from simple job bundles that users submit to a queue with either the CLI or an automatically generated GUI, to integrated submitter plugins that dynamically generate a job bundle for an application-defined workload.

How workloads arise from production

To understand workloads in production contexts and how to support them with Deadline Cloud, consider how they come to be. Production may involve creating visual effects, animation, games, product catalog imagery, 3D reconstructions for building information modeling (BIM), and more. This content is typically created by a team of artistic or technical specialists running a variety of software applications and custom scripting. Members of the team pass data between each other using a production pipeline. Many tasks performed by the pipeline involve intensive computations that would take days if run on a user's workstation.

Some examples of tasks in these production pipelines include:

- Using a photogrammetry application to process photographs taken of a movie set to reconstruct a textured digital mesh.
- Running a particle simulation in a 3D scene to add layers of detail to an explosion visual effect for a television show.
- Cooking data for a game level into the form needed for external release and applying optimization and compression settings.
- Rendering a set of images for a product catalog including variations in color, background, and lighting.
- Running a custom-developed script on a 3D model to apply a look that was custom-built and approved by a movie director.

These tasks involve many parameters to adjust to get an artistic result or to fine tune the output quality. Often there is a GUI to select those parameter values with a button or menu to run

the process locally within the application. When a user runs the process, the application and possibly the host computer itself cannot be used to perform other operations because it uses the application state in memory and may consume all of the host computer's CPU and memory resources.

In many cases the process is quick. During the course of production, the speed of the process slows down when the requirements for quality and complexity go up. A character test that took 30 seconds during development can easily turn into 3 hours when it is applied to the final production character. Through this progression, a workload that began life inside a GUI can grow too large to fit. Porting it to Deadline Cloud can boost the productivity of users running these processes because they get back full control of their workstation and can keep track of more iterations from the Deadline Cloud monitor.

There are two levels of support to aim for when developing support for a workload in Deadline Cloud:

- Offloading the workload from the user workstation to a Deadline Cloud farm with no parallelism or speed-up. This may under-utilize the available compute resources in the farm, but the ability to shift long operations to a batch processing system enables users to get more done with their own workstation.
- Optimizing the parallelism of the workload so that it utilizes the Deadline Cloud farm's horizontal scale to complete quickly.

There are times that it is obvious how to make a workload run in parallel. For example, each frame of a computer graphics render can be done independently. It's important not to get stuck on this parallelism, however. Instead, understand that offloading a long-running workload to Deadline Cloud provides significant benefits, even when there is no obvious way to split the workload up.

The ingredients of a workload

To specify a Deadline Cloud workload, implement a job bundle that users submit to a queue with the [Deadline Cloud CLI](). Much of the work in creating a job bundle is to write the job template, but there are more factors like how to provide the applications that the workload requires. Here are the essential things to consider when defining a workload for Deadline Cloud:

- **The application to run**. The job must be able to launch application processes, and therefore needs an installation of the application available as well as any licensing the application uses,

such as access to a floating license server. This is typically part of the farm configuration, and not embedded in the job bundle itself.

- Configure jobs using queue environments
- Connect customer-managed fleets to a license endpoint
- **Job parameter definitions**. The user experience of submitting the job is affected greatly by the parameters it provides. Example parameters include data files, directories, and application configuration.
  - Parameter values elements for job bundles
- **File data flow**. When a job runs, it reads input from files provided by the user, then writes its output as new files. To work with the job attachments and path mapping features, the job must specify the paths of the directories or specific files for these inputs and outputs.
  - Using files in your jobs
- **The step script**. The step script runs the application binary with the right command-line options to apply the provided job parameters. It also handles details like path mapping if the workload data files include absolute instead of relative path references.
  - Job template elements for job bundles

Workload portability

A workload is portable when it can run in multiple different systems without changing it each time you submit a job. For example, it might run on different render farms that have different shared file systems mounted, or on different operating systems like Linux or Windows. When you implement a portable job bundle, it's easier for users to run the job on their specific farm, or to adapt it for other use cases.

Here are some ways you can make your job bundle portable.

- Fully specify the input data files needed by a workload, using PATH job parameters and asset references in the job bundle. This makes the job portable to farms based on shared file systems and to farms that make copies of the input data, like the Deadline Cloud job attachments feature.
- Make file path references for the input files of the job relocatable and usable on different operating systems. For example when users submit jobs from Windows workstations to run on a Linux fleet.
  - Use relative file path references, so if the directory containing them is moved to a different location, references still resolve. Some applications, like Blender, support a choice between relative and absolute paths.

- If you can't use relative paths, support OpenJD [path mapping metadata](#) and translate the absolute paths according to how Deadline Cloud provides the files to the job.
- Implement commands in a job using portable scripts. Python and bash are two examples of scripting languages that can be used this way. You should consider providing them both on all the worker hosts of your fleets.
  - Use the script interpreter binary, like `python` or `bash`, with the script file name as an argument. This works on all operating systems including Windows, compared to using a script file with its execute bit set on Linux.
  - Write portable bash scripts by applying these practices:
    - Expand template path parameters in single quotes to handle paths with spaces and Windows path separators.
    - When running on Windows, watch for issues related to MinGW automatic path translation. For example, it transforms an AWS CLI command like `aws logs tail /aws/ deadline/...` into a command similar to `aws logs tail "C:/Program Files/Git/ aws/deadline/..."` and won't tail a log correctly. Set the variable `MSYS_NO_PATHCONV=1` to turn this behavior off.
    - In most cases, the same code works on all operating systems. When the code needs to be different use an `if/else` construct to handle the cases.

      ```
      if [[ "$(uname)" == MINGW* ]]; then
          # Code for Windows
      elif [[ "$(uname)" == Darwin ]]; then
          # Code for MacOS
      else
          # Code for Linux and other operating systems
      fi
      ```

  - You can write portable Python scripts using `pathlib` to handle file system path differences and avoid operating-specific features. The Python documentation includes annotations for this, for example in the [signal library documentation](#). Linux-specific feature support is marked as "Availability: Linux."
- Use job parameters to specify application requirements. Use consistent conventions that the farm administrator can apply in [queue environments](#).
  - For example, you can use the `CondaPackages` and/or `RezPackages` parameters in your job, with a default parameter value that lists the application package names and versions the job requires. Then, you can use one of the [sample Conda or Rez queue environments](#) to provide a virtual environment for the job.

Getting started with Deadline Cloud resources.

To start creating custom solutions for AWS Deadline Cloud, you must set up your resources. These include a farm, at least one queue for the farm, and at least one worker fleet to service the queue. You can create your resources using the Deadline Cloud console, or you can use the AWS Command Line Interface.

In this tutorial, you will use AWS CloudShell to create a simple developer farm and run the worker agent. You can then submit and run a simple job with parameters and attachments, add a service managed fleet, and clean up your farm resources when you're done.

The following sections introduce you to the different features of Deadline Cloud, and how they function and work together. Following these steps is useful for developing and testing new workloads and customizations.

For instructions to set up your farm using the console, see Getting started in the *Deadline Cloud User Guide*.

Topics

- Create a Deadline Cloud farm
- Run the Deadline Cloud worker agent
- Submit with Deadline Cloud
- Submit jobs with job attachments in Deadline Cloud
- Add a service-managed fleet to your developer farm in Deadline Cloud
- Clean up your farm resources in Deadline Cloud

Create a Deadline Cloud farm

To create your developer farm and queue resources in AWS Deadline Cloud, use the AWS Command Line Interface (AWS CLI), as shown in the following procedure. You will also create an AWS Identity and Access Management (IAM) role and a customer-managed fleet (CMF) and associate the fleet with your queue. Then you can configure the AWS CLI and confirm that your farm is set up and working as specified.

You can use this farm to explore the features of Deadline Cloud, then develop and test new workloads, customizations, and pipeline integrations.

To create a farm

1. [Open an AWS CloudShell session](). You'll use the CloudShell window to enter AWS Command Line Interface (AWS CLI) commands to run the examples in this tutorial. Keep the CloudShell window open as you proceed.

2. Create a name for your farm, and add that farm name to `~/.bashrc`. This will make it available for other terminal sessions.

```
echo "DEV_FARM_NAME=DeveloperFarm" >> ~/.bashrc
source ~/.bashrc
```

3. Create the farm resource, and add its farm ID to `~/.bashrc`.

```
aws deadline create-farm \
    --display-name "$DEV_FARM_NAME"

echo "DEV_FARM_ID=\$(aws deadline list-farms \
        --query \"farms[?displayName=='\$DEV_FARM_NAME'].farmId \
        | [0]\" --output text)" >> ~/.bashrc
source ~/.bashrc
```

4. Create the queue resource, and add its queue ID to `~/.bashrc.`

```
aws deadline create-queue \
    --farm-id $DEV_FARM_ID \
    --display-name "$DEV_FARM_NAME Queue" \
    --job-run-as-user '{"posix": {"user": "job-user", "group": "job-group"},
 "runAs":"QUEUE_CONFIGURED_USER"}'

echo "DEV_QUEUE_ID=\$(aws deadline list-queues \
        --farm-id \$DEV_FARM_ID \
        --query \"queues[?displayName=='\$DEV_FARM_NAME Queue'].queueId \
        | [0]\" --output text)" >> ~/.bashrc
source ~/.bashrc
```

5. Create an IAM role for the fleet. This role provides worker hosts in your fleet with the necessary security credentials to run jobs from your queue.

```
aws iam create-role \
    --role-name "${DEV_FARM_NAME}FleetRole" \
    --assume-role-policy-document \
        '{
```

```
                "Version": "2012-10-17",
                "Statement": [
                    {
                        "Effect": "Allow",
                        "Principal": {
                            "Service": "credentials.deadline.amazonaws.com"
                        },
                        "Action": "sts:AssumeRole"
                    }
                ]
            }'
aws iam put-role-policy \
    --role-name "${DEV_FARM_NAME}FleetRole" \
    --policy-name WorkerPermissions \
    --policy-document \
        '{
            "Version": "2012-10-17",
            "Statement": [
                {
                    "Effect": "Allow",
                    "Action": [
                        "deadline:AssumeFleetRoleForWorker",
                        "deadline:UpdateWorker",
                        "deadline:DeleteWorker",
                        "deadline:UpdateWorkerSchedule",
                        "deadline:BatchGetJobEntity",
                        "deadline:AssumeQueueRoleForWorker"
                    ],
                    "Resource": "*",
                    "Condition": {
                        "StringEquals": {
                            "aws:PrincipalAccount": "${aws:ResourceAccount}"
                        }
                    }
                },
                {
                    "Effect": "Allow",
                    "Action": [
                        "logs:CreateLogStream"
                    ],
                    "Resource": "arn:aws:logs:*:*:*:/aws/deadline/*",
                    "Condition": {
                        "StringEquals": {
                            "aws:PrincipalAccount": "${aws:ResourceAccount}"
```

```
                } 
            } 
        }, 
        { 
            "Effect": "Allow", 
            "Action": [ 
                "logs:PutLogEvents", 
                "logs:GetLogEvents" 
            ], 
            "Resource": "arn:aws:logs:*:*:*:/aws/deadline/*", 
            "Condition": { 
                "StringEquals": { 
                    "aws:PrincipalAccount": "${aws:ResourceAccount}" 
                } 
            } 
        } 
    ] 
}'
```

6. Create the customer-managed fleet (CMF), and add its fleet ID to ~/.bashrc.

```
FLEET_ROLE_ARN="arn:aws:iam::$(aws sts get-caller-identity \ 
        --query "Account" --output text):role/${DEV_FARM_NAME}FleetRole" 
aws deadline create-fleet \ 
    --farm-id $DEV_FARM_ID \ 
    --display-name "$DEV_FARM_NAME CMF" \ 
    --role-arn $FLEET_ROLE_ARN \ 
    --max-worker-count 5 \ 
    --configuration \ 
        '{ 
            "customerManaged": { 
                "mode": "NO_SCALING", 
                "workerCapabilities": { 
                    "vCpuCount": {"min": 1}, 
                    "memoryMiB": {"min": 512}, 
                    "osFamily": "linux", 
                    "cpuArchitectureType": "x86_64" 
                } 
            } 
        }' 

echo "DEV_CMF_ID=\$(aws deadline list-fleets \ 
        --farm-id \$DEV_FARM_ID \ 
        --query \"fleets[?displayName=='\$DEV_FARM_NAME CMF'].fleetId \
```

```
        | [0]\" --output text)" >> ~/.bashrc
source ~/.bashrc
```

7. Ensure you can access Deadline Cloud.

```
pip install deadline
```

8. Associate the CMF with your queue.

```
aws deadline create-queue-fleet-association \
    --farm-id $DEV_FARM_ID \
    --queue-id $DEV_QUEUE_ID \
    --fleet-id $DEV_CMF_ID
```

9. To set the default farm to the farm ID and the queue to the queue ID that you created earlier, use the following command.

```
deadline config set defaults.farm_id $DEV_FARM_ID
deadline config set defaults.queue_id $DEV_QUEUE_ID
```

10. (Optional) To confirm that your farm is set up according to your specifications, use the following commands:

- List all farms – **deadline farm list**
- List all queues in the default farm – **deadline queue list**
- List all fleets in the default farm – **deadline fleet list**
- Get the default farm – **deadline farm get**
- Get the default queue – **deadline queue get**
- Get all the fleets associated with the default queue – **deadline fleet get**

Next steps

After you create your farm, you can run the Deadline Cloud worker agent on the hosts in your fleet to process jobs. See [Run the Deadline Cloud worker agent](#).

Run the Deadline Cloud worker agent

Before you can run the jobs you submit to the queue on your developer farm, you must run the AWS Deadline Cloud worker agent in developer mode on a worker host.

Throughout the remainder of this tutorial, you will perform AWS CLI operations on your developer farm using two AWS CloudShell tabs. In the first tab, you can submit jobs. In the second tab, you can run the worker agent.

> **ⓘ Note**
>
> If you leave your CloudShell session idle for more than 20 minutes, it will timeout and stop the worker agent. To restart the worker agent, follow the instructions in the following procedure.

Before you can start a worker agent, you must set up a Deadline Cloud farm, queue, and fleet. See [Create a Deadline Cloud farm](#).

To run the worker agent in developer mode

1. With your farm still open in the first CloudShell tab, open a second CloudShell tab, then create the `demoenv-logs` and `demoenv-persist` directories.

   ```
   mkdir ~/demoenv-logs
   mkdir ~/demoenv-persist
   ```

2. Download and install the Deadline Cloud worker agent packages from PyPI:

   > **ⓘ Note**
   >
   > On Windows, it is required that the agent files are installed into Python's global site-packages directory. Python virtual environments are not currently supported.

   ```
   python -m pip install deadline-cloud-worker-agent
   ```

3. To allow the worker agent to create the temporary directories for running jobs, create a directory:

   ```
   sudo mkdir /sessions
   sudo chmod 750 /sessions
   sudo chown cloudshell-user /sessions
   ```

4. Run the Deadline Cloud worker agent in developer mode with the variables DEV_FARM_ID and DEV_CMF_ID that you added to the ~/.bashrc.

```
deadline-worker-agent \
    --farm-id $DEV_FARM_ID \
    --fleet-id $DEV_CMF_ID \
    --run-jobs-as-agent-user \
    --logs-dir ~/demoenv-logs \
    --persistence-dir ~/demoenv-persist
```

As the worker agent initializes and then polls the UpdateWorkerSchedule API operation the following output is displayed:

```
INFO     Worker Agent starting
[2024-03-27 15:51:01,292][INFO     ] # Worker Agent starting
[2024-03-27 15:51:01,292][INFO     ] AgentInfo
Python Interpreter: /usr/bin/python3
Python Version: 3.9.16 (main, Sep  8 2023, 00:00:00)  - [GCC 11.4.1 20230605 (Red
 Hat 11.4.1-2)]
Platform: linux
...
[2024-03-27 15:51:02,528][INFO     ] # API.Resp # [deadline:UpdateWorkerSchedule]
(200) params={'assignedSessions': {}, 'cancelSessionActions': {},
 'updateIntervalSeconds': 15} ...
[2024-03-27 15:51:17,635][INFO     ] # API.Resp # [deadline:UpdateWorkerSchedule]
(200) params=(Duplicate removed, see previous response) ...
[2024-03-27 15:51:32,756][INFO     ] # API.Resp # [deadline:UpdateWorkerSchedule]
(200) params=(Duplicate removed, see previous response) ...
...
```

5. Select your first CloudShell tab, then list the workers in the fleet.

```
deadline worker list --fleet-id $DEV_CMF_ID
```

Output such as the following is displayed:

```
Displaying 1 of 1 workers starting at 0

- workerId: worker-8c9af877c8734e89914047111f
  status: STARTED
  createdAt: 2023-12-13 20:43:06+00:00
```

In a production configuration, the Deadline Cloud worker agent requires setting up multiple users and configuration directories as an administrative user on the host machine. You can override these settings because you're running jobs in your own development farm, which only you can access.

Next steps

Now that a worker agent is running on your worker hosts, you can send jobs to your workers. You can:

- Submit with Deadline Cloud using a simple OpenJD job bundle.

- Submit jobs with job attachments in Deadline Cloud that share files between workstations using different operating systems.

Submit with Deadline Cloud

To run Deadline Cloud jobs on your worker hosts, you create and use an Open Job Description (OpenJD) job bundle to configure a job. The bundle configures the job, for example by specifying input files for a job and where to write the output of the job. This topic includes examples of ways that you can configure a job bundle.

Before you can follow the procedures in this section, you must complete the following:

- Create a Deadline Cloud farm

- Run the Deadline Cloud worker agent

To use AWS Deadline Cloud to run jobs, use the following procedures. Use the first AWS CloudShell tab to submit jobs to your developer farm. Use the second CloudShell tab to view the worker agent output.

Topics

- Submit the simple_job sample

- Submit a simple_job with a parameter

- Create a simple_file_job job bundle with file I/O

- Next steps

Submit the simple_job sample

After you create a farm and run the worker agent, you can submit the simple_job sample to Deadline Cloud.

To submit the simple_job sample to Deadline Cloud

1. Download the sample from GitHub.

    ```
    cd ~
    git clone https://github.com/aws-deadline/deadline-cloud-samples.git
    ```

2. Choose your first CloudShell tab, then navigate to the job bundle samples directory.

    ```
    cd ~/deadline-cloud-samples/job_bundles/
    ```

3. Submit the simple_job sample.

    ```
    deadline bundle submit simple_job
    ```

4. Choose your second CloudShell tab to view the logging output about calling `BatchGetJobEntities`, getting a session, and running a session action.

    ```
    ...
    [2024-03-27 16:00:21,846][INFO    ] # Session.Starting
     # [session-053d77cef82648fe2] Starting new Session.
     [queue-3ba4ff683ff54db09b851a2ed8327d7b/job-d34cc98a6e234b6f82577940ab4f76c6]
    [2024-03-27 16:00:21,853][INFO    ] # API.Req # [deadline:BatchGetJobEntity]
     resource={'farm-id': 'farm-3e24cfc9bbcd423e9c1b6754bc1',
     'fleet-id': 'fleet-246ee60f46d44559b6cce010d05', 'worker-id':
     'worker-75e0fce9c3c344a69bff57fcd83'} params={'identifiers': [{'jobDetails':
     {'jobId': 'job-d34cc98a6e234b6f82577940ab4'}}]} request_url=https://
    scheduling.deadline.us-west-2.amazonaws.com/2023-10-12/farms/
    farm-3e24cfc9bbcd423e /fleets/fleet-246ee60f46d44559b1 /workers/worker-
     75e0fce9c3c344a69b /batchGetJobEntity
    [2024-03-27 16:00:22,013][INFO    ] # API.Resp # [deadline:BatchGetJobEntity](200)
     params={'entities': [{'jobDetails': {'jobId': 'job-d34cc98a6e234b6f82577940ab6',
     'jobRunAsUser': {'posix': {'user': 'job-user', 'group': 'job-group'},
     'runAs': 'QUEUE_CONFIGURED_USER'}, 'logGroupName': '/aws/deadline/
    farm-3e24cfc9bbcd423e9c1b6754bc1/queue-3ba4ff683ff54db09b851a2ed83', 'parameters':
     '*REDACTED*', 'schemaVersion': 'jobtemplate-2023-09'}}], 'errors': []}
     request_id=a3f55914-6470-439e-89e5-313f0c6
    ```

```
[2024-03-27 16:00:22,013][INFO     ] # Session.Add #
 [session-053d77cef82648fea9c69827182] Appended new SessionActions.
 (ActionIds: ['sessionaction-053d77cef82648fea9c69827182-0'])
 [queue-3ba4ff683ff54db09b851a2ed8b/job-d34cc98a6e234b6f82577940ab6]
[2024-03-27 16:00:22,014][WARNING ] # Session.User #
 [session-053d77cef82648fea9c69827182] Running as the Worker Agent's
 user. (User: cloudshell-user) [queue-3ba4ff683ff54db09b851a2ed8b/job-
 d34cc98a6e234b6f82577940ac6]
[2024-03-27 16:00:22,015][WARNING ] # Session.AWSCreds #
 [session-053d77cef82648fea9c69827182] AWS Credentials are not available: Queue has
 no IAM Role. [queue-3ba4ff683ff54db09b851a2ed8b/job-d34cc98a6e234b6f82577940ab6]
[2024-03-27 16:00:22,026][INFO     ] # Session.Logs #
 [session-053d77cef82648fea9c69827182] Logs streamed to: AWS CloudWatch
 Logs. (LogDestination: /aws/deadline/farm-3e24cfc9bbcd423e9c1b6754bc1/
 queue-3ba4ff683ff54db09b851a2ed83/session-053d77cef82648fea9c69827181)
 [queue-3ba4ff683ff54db09b851a2ed83/job-d34cc98a6e234b6f82577940ab4]
[2024-03-27 16:00:22,026][INFO     ] # Session.Logs #
 [session-053d77cef82648fea9c69827182] Logs streamed to: local
 file. (LogDestination: /home/cloudshell-user/demoenv-logs/
 queue-3ba4ff683ff54db09b851a2ed8b/session-053d77cef82648fea9c69827182.log)
 [queue-3ba4ff683ff54db09b851a2ed83/job-d34cc98a6e234b6f82577940ab4]
...
```

> ⓘ **Note**
>
> Only the logging output from the worker agent is shown. There is a separate log for
> the session that runs the job.

5. Choose your first tab, then inspect the log files that the worker agent writes.

    a. Navigate to the worker agent logs directory and view its contents.

    ```
    cd ~/demoenv-logs
    ls
    ```

    b. Print the first log file that the worker agent creates.

    ```
    cat worker-agent-bootstrap.log
    ```

    This file contains worker agent output about how it called the Deadline Cloud API to
    create a worker resource in your fleet, and then assumed the fleet role.

c.   Print the log file output when the worker agent joins the fleet.

```
cat worker-agent.log
```

This log contains outputs about all the actions that the worker agent takes, but doesn't contain output about the queues it runs jobs from, except for the IDs of those resources.

d.   Print the log files for each session in a directory that is named the same as the queue resource id.

```
cat $DEV_QUEUE_ID/session-*.log
```

If the job is successful, the log file output will be similar to the following:

```
cat $DEV_QUEUE_ID/$(ls -t $DEV_QUEUE_ID | head -1)

2024-03-27 16:00:22,026 WARNING Session running with no AWS Credentials.
2024-03-27 16:00:22,404 INFO
2024-03-27 16:00:22,405 INFO ==============================================
2024-03-27 16:00:22,405 INFO --------- Running Task
2024-03-27 16:00:22,405 INFO ==============================================
2024-03-27 16:00:22,406 INFO ----------------------------------------------
2024-03-27 16:00:22,406 INFO Phase: Setup
2024-03-27 16:00:22,406 INFO ----------------------------------------------
2024-03-27 16:00:22,406 INFO Writing embedded files for Task to disk.
2024-03-27 16:00:22,406 INFO Mapping: Task.File.runScript -> /sessions/
session-053d77cef82648fea9c698271812a/embedded_fileswa_gj55_/tmp2u9yqtsz
2024-03-27 16:00:22,406 INFO Wrote: runScript -> /sessions/
session-053d77cef82648fea9c698271812a/embedded_fileswa_gj55_/tmp2u9yqtsz
2024-03-27 16:00:22,407 INFO ----------------------------------------------
2024-03-27 16:00:22,407 INFO Phase: Running action
2024-03-27 16:00:22,407 INFO ----------------------------------------------
2024-03-27 16:00:22,407 INFO Running command /sessions/
session-053d77cef82648fea9c698271812a/tmpzuzxpslm.sh
2024-03-27 16:00:22,414 INFO Command started as pid: 471
2024-03-27 16:00:22,415 INFO Output:
2024-03-27 16:00:22,420 INFO Welcome to AWS Deadline Cloud!
2024-03-27 16:00:22,571 INFO
2024-03-27 16:00:22,572 INFO ==============================================
2024-03-27 16:00:22,572 INFO --------- Session Cleanup
2024-03-27 16:00:22,572 INFO ==============================================
```

```
2024-03-27 16:00:22,572 INFO Deleting working directory: /sessions/
session-053d77cef82648fea9c698271812a
```

6. Print information about the job.

```
deadline job get
```

When you submit the job, the system saves it as the default so you don't have to enter the job ID.

Submit a simple_job with a parameter

You can submit jobs with parameters. In the following procedure, you edit the simple_job template to include a custom message, submit the simple_job, then print the session log file to view the message.

To submit the simple_job sample with a parameter

1. Select your first CloudShell tab, then navigate to the job bundle samples directory.

```
cd ~/deadline-cloud-samples/job_bundles/
```

2. Print the contents of the simple_job template.

```
cat simple_job/template.yaml
```

The `parameterDefinitions` section with the `Message` parameter should look like the following:

```
parameterDefinitions:
- name: Message
  type: STRING
  default: Welcome to AWS Deadline Cloud!
```

3. Submit the simple_job sample with a parameter value, then wait for the job to finish running.

```
deadline bundle submit simple_job \
    -p "Message=Greetings from the developer getting started guide."
```

4. To see the custom message, view the most recent session log file.

```
cd ~/demoenv-logs
cat $DEV_QUEUE_ID/$(ls -t $DEV_QUEUE_ID | head -1)
```

Create a simple_file_job job bundle with file I/O

A render job needs to read the scene definition, render an image from it, and then save that image to an output file. You can simulate this action by making the job compute the hash of the input instead of rendering an image.

To create a simple_file_job job bundle with file I/O

1. Select your first CloudShell tab, then navigate to the job bundle samples directory.

   ```
   cd ~/deadline-cloud-samples/job_bundles/
   ```

2. Make a copy of `simple_job` with the new name `simple_file_job`.

   ```
   cp -r simple_job simple_file_job
   ```

3. Edit the job template as follows:

   > **ⓘ Note**
   >
   > We recommend that you use nano for these steps. If you prefer to use Vim, you must set its paste mode using `:set paste`.

   a. Open the template in a text editor.

   ```
   nano simple_file_job/template.yaml
   ```

   b. Add the following `type`, `objectType`, and `dataFlow parameterDefinitions`.

   ```
   - name: InFile
     type: PATH
     objectType: FILE
     dataFlow: IN
   - name: OutFile
     type: PATH
   ```

```
    objectType: FILE
    dataFlow: OUT
```

c. Add the following `bash` script command to the end of the file that reads from the input file and writes to the output file.

```
# hash the input file, and write that to the output
sha256sum "{{Param.InFile}}" > "{{Param.OutFile}}"
```

The updated `template.yaml` should exactly match the following:

```
specificationVersion: 'jobtemplate-2023-09'
name: Simple File Job Bundle Example
parameterDefinitions:
- name: Message
  type: STRING
  default: Welcome to AWS Deadline Cloud!
- name: InFile
  type: PATH
  objectType: FILE
  dataFlow: IN
- name: OutFile
  type: PATH
  objectType: FILE
  dataFlow: OUT
steps:
- name: WelcomeToDeadlineCloud
  script:
    actions:
      onRun:
        command: '{{Task.File.runScript}}'
    embeddedFiles:
    - name: runScript
      type: TEXT
      runnable: true
      data: |
        #!/usr/bin/env bash
        echo "{{Param.Message}}"

        # hash the input file, and write that to the output
        sha256sum "{{Param.InFile}}" > "{{Param.OutFile}}"
```

> **ⓘ Note**
>
> If you want to adjust the spacing in the `template.yaml`, make sure that you use spaces instead of indentations.

    d.   Save the file, and exit the text editor.

4.   Provide parameter values for the input and output files to submit the simple_file_job.

```
deadline bundle submit simple_file_job \
    -p "InFile=simple_job/template.yaml" \
    -p "OutFile=hash.txt"
```

5.   Print information about the job.

```
deadline job get
```

- You will see output such as the following:

```
parameters:
  Message:
    string: Welcome to AWS Deadline Cloud!
  InFile:
    path: /local/home/cloudshell-user/BundleFiles/JobBundle-Examples/simple_job/
template.yaml
  OutFile:
    path: /local/home/cloudshell-user/BundleFiles/JobBundle-Examples/hash.txt
```

- Although you only provided relative paths, the parameters have the full path set. The AWS CLI joins the current working directory to any paths that are provided as parameters when the paths have the type PATH.

- The worker agent running in the other terminal window picks up and runs the job. This action creates the `hash.txt` file, which you can view with the following command.

```
cat hash.txt
```

This command will print output similar to the following.

```
eaa2df5d34b54be5ac34c56a24a8c237b8487231a607eaf530a04d76b89c9cd3 /local/home/
cloudshell-user/BundleFiles/JobBundle-Examples/simple_job/template.yaml
```

Next steps

After learning how to submit simple jobs using the Deadline Cloud CLI, you can explore:

- [Submit jobs with job attachments in Deadline Cloud](#) to learn how to run jobs on hosts running different operating systems.

- [Add a service-managed fleet to your developer farm in Deadline Cloud](#) to run your jobs on hosts managed by Deadline Cloud.

- [Clean up your farm resources in Deadline Cloud](#) to shut down the resources that you used for this tutorial.

Submit jobs with job attachments in Deadline Cloud

Many farms use shared filesystems to share files between the hosts that submit jobs and those that run jobs. For example, in the previous `simple_file_job` example, the local filesystem is shared between the AWS CloudShell terminal windows, which run in tab one where you submit the job, and tab two where you run the worker agent.

A shared filesystem is advantageous when the submitter workstation and the worker hosts are on the same local area network. If you store your data on premises near the workstations that access it, then using a cloud-based farm means you have to share your filesystems over a high-latency VPN or synchronize your filesystems in the cloud. Neither of these options are easy to set up or operate.

AWS Deadline Cloud offers a simple solution with *job attachments*, which are similar to email attachments. With job attachments, you attach data to your job. Then, Deadline Cloud handles the details of transferring and storing your job data in Amazon Simple Storage Service (Amazon S3) buckets.

Content creation workflows are often iterative, meaning a user submits jobs with a small subset of modified files. Because Amazon S3 buckets store job attachments in a content-addressable storage, the name of each object is based on the hash of the object's data and the contents of a directory tree are stored in a manifest file format attached to a job.

Before you can follow the procedures in this section, you must complete the following:

- [Create a Deadline Cloud farm](#)

- [Run the Deadline Cloud worker agent](#)

To run jobs with job attachments, complete the following steps.

Topics

- [Add a job attachments configuration to your queue](#)

- [Submit simple_file_job with job attachments](#)

- [Understanding how job attachments are stored in Amazon S3](#)

- [Next steps](#)

Add a job attachments configuration to your queue

To enable job attachments in your queue, add a job attachments configuration to the queue resource in your account.

To add a job attachments configuration to your queue

1. Choose your first CloudShell tab, then enter one of the following commands to use an Amazon S3 bucket for job attachments.

   - If you don't have an existing private Amazon S3 bucket, you can create and use a new S3 bucket.

     ```
     DEV_FARM_BUCKET=$(echo $DEV_FARM_NAME \
         | tr '[:upper:]' '[:lower:]')-$(xxd -l 16 -p /dev/urandom)
     if [ "$AWS_REGION" == "us-east-1" ]; then LOCATION_CONSTRAINT=
     else LOCATION_CONSTRAINT="--create-bucket-configuration \
         LocationConstraint=${AWS_REGION}"
     fi
     aws s3api create-bucket \
         $LOCATION_CONSTRAINT \
         --acl private \
         --bucket ${DEV_FARM_BUCKET}
     ```

   - If you already have a private Amazon S3 bucket, you can use it by replacing *MY_BUCKET_NAME* with the name of your bucket.

```
DEV_FARM_BUCKET=MY_BUCKET_NAME
```

2. After you create or choose your Amazon S3 bucket, add the bucket name to `~/.bashrc` to make the bucket available for other terminal sessions.

```
echo "DEV_FARM_BUCKET=$DEV_FARM_BUCKET" >> ~/.bashrc
```

3. Create an AWS Identity and Access Management (IAM) role for the queue.

```
aws iam create-role --role-name "${DEV_FARM_NAME}QueueRole" \
    --assume-role-policy-document \
        '{
            "Version": "2012-10-17",
            "Statement": [
                {
                    "Effect": "Allow",
                    "Principal": {
                        "Service": "credentials.deadline.amazonaws.com"
                    },
                    "Action": "sts:AssumeRole"
                }
            ]
        }'
aws iam put-role-policy \
    --role-name "${DEV_FARM_NAME}QueueRole" \
    --policy-name S3BucketsAccess \
    --policy-document \
        '{
            "Version": "2012-10-17",
            "Statement": [
                {
                    "Action": [
                        "s3:GetObject*",
                        "s3:GetBucket*",
                        "s3:List*",
                        "s3:DeleteObject*",
                        "s3:PutObject",
                        "s3:PutObjectLegalHold",
                        "s3:PutObjectRetention",
                        "s3:PutObjectTagging",
                        "s3:PutObjectVersionTagging",
                        "s3:Abort*"
```

```
                    ],
                    "Resource": [
                        "arn:aws:s3:::'$DEV_FARM_BUCKET'",
                        "arn:aws:s3:::'$DEV_FARM_BUCKET'/*"
                    ],
                    "Effect": "Allow"
                }
            ]
            }'
```

4. Update your queue to include the job attachments settings and the IAM role.

```
QUEUE_ROLE_ARN="arn:aws:iam::$(aws sts get-caller-identity \
        --query "Account" --output text):role/${DEV_FARM_NAME}QueueRole"
aws deadline update-queue \
    --farm-id $DEV_FARM_ID \
    --queue-id $DEV_QUEUE_ID \
    --role-arn $QUEUE_ROLE_ARN \
    --job-attachment-settings \
        '{
            "s3BucketName": "'$DEV_FARM_BUCKET'",
            "rootPrefix": "JobAttachments"
        }'
```

5. Confirm that you updated your queue.

```
deadline queue get
```

Output such as the following is shown:

```
...
jobAttachmentSettings:
  s3BucketName: DEV_FARM_BUCKET
  rootPrefix: JobAttachments
roleArn: arn:aws:iam::ACCOUNT_NUMBER:role/DeveloperFarmQueueRole
...
```

Submit simple_file_job with job attachments

When you use job attachments, job bundles must give Deadline Cloud enough information to determine the job's data flow, such as using PATH parameters. In the case of the simple_file_job,

you edited the `template.yaml` file to tell Deadline Cloud that the data flow is in the input file and output file.

After you've added the job attachments configuration to your queue, you can submit the simple_file_job sample with job attachments. After you do this, you can view the logging and job output to confirm that the simple_file_job with job attachments is working.

To submit the simple_file_job job bundle with job attachments

1. Choose your first CloudShell tab, then open the `JobBundle-Samples` directory.

2.
```
cd ~/AmazonDeadlineCloud-DocumentationAndSamples/JobBundle-Samples
```

3. Submit simple_file_job to the queue. When prompted to confirm the upload, enter **y**.

```
deadline bundle submit simple_file_job \
    -p InFile=simple_job/template.yaml \
    -p OutFile=hash-jobattachments.txt
```

4. To view the job attachments data transfer session log output, choose your second CloudShell tab.

```
JOB_ID=$(deadline config get defaults.job_id)
SESSION_ID=$(aws deadline list-sessions \
        --farm-id $DEV_FARM_ID \
        --queue-id $DEV_QUEUE_ID \
        --job-id $JOB_ID \
        --query "sessions[0].sessionId" \
        --output text)
cat ~/demoenv-logs/$DEV_QUEUE_ID/$$SESSION_ID.log
```

5. List the session actions that were run within the session.

```
aws deadline list-session-actions \
    --farm-id $DEV_FARM_ID \
    --queue-id $DEV_QUEUE_ID \
    --job-id $JOB_ID \
    --session-id $SESSION_ID
```

Output such as the following is shown:

```
{
```

```
    "sessionactions": [
        {
            "sessionActionId": "sessionaction-123-0",
            "status": "SUCCEEDED",
            "startedAt": "<timestamp>",
            "endedAt": "<timestamp>",
            "progressPercent": 100.0,
            "definition": {
                "syncInputJobAttachments": {}
            }
        },
        {
            "sessionActionId": "sessionaction-123-1",
            "status": "SUCCEEDED",
            "startedAt": "<timestamp>",
            "endedAt": "<timestamp>",
            "progressPercent": 100.0,
            "definition": {
                "taskRun": {
                    "taskId": "task-abc-0",
                    "stepId": "step-def"
                }
            }
        }
    ]
}
```

The first session action downloaded the input job attachments, while the second action runs the task like before and then uploaded the output job attachments.

6. List the output directory.

```
ls *.txt
```

Output such as hash.txt is shown, but hash-jobattachments.txt doesn't exist.

7. Download the output from the most recent job.

```
deadline job download-output
```

8. View the output of the downloaded file.

```
cat hash-jobattachments.txt
```

Output such as the following is shown:

```
eaa2df5d34b54be5ac34c56a24a8c237b8487231a607eaf530a04d76b89c9cd3   /tmp/openjd/
session-123/assetroot-abc/simple_job/template.yaml
```

Understanding how job attachments are stored in Amazon S3

You can use the AWS Command Line Interface (AWS CLI) to upload or download data for job attachments, which are stored in Amazon S3 buckets. Understanding how Deadline Cloud stores job attachments on Amazon S3 will help when you develop workloads and pipeline integrations.

To inspect how Deadline Cloud job attachments are stored in Amazon S3

1. Choose your first CloudShell tab, then open the job bundle samples directory.

    ```
    cd ~/AmazonDeadlineCloud-DocumentationAndSamples/JobBundle-Samples
    ```

2. Inspect the job properties.

    ```
    deadline job get
    ```

    Output such as the following is shown:

    ```
    parameters:
      Message:
        string: Welcome to Amazon Deadline Cloud!
      InFile:
        path: /home/cloudshell-user/AmazonDeadlineCloud-DocumentationAndSamples/
    JobBundle-Samples/simple_job/template.yaml
      OutFile:
        path: /home/cloudshell-user/AmazonDeadlineCloud-DocumentationAndSamples/
    JobBundle-Samples/hash-jobattachments.txt
    attachments:
      manifests:
      - rootPath: /home/cloudshell-user/AmazonDeadlineCloud-DocumentationAndSamples/
    JobBundle-Samples
        rootPathFormat: posix
    ```

```
    outputRelativeDirectories:
    - .
    inputManifestPath: farm-3040c59a5b9943d58052c29d907a645d/queue-
  cde9977c9f4d4018a1d85f3e6c1a4e6e/Inputs/
  f46af01ca8904cd8b514586671c79303/0d69cd94523ba617c731f29c019d16e8_input.xxh128
    inputManifestHash: f95ef91b5dab1fc1341b75637fe987ee
  fileSystem: COPIED
```

The attachments field contains a list of manifest structures that describe input and output data paths that the job uses when it runs. Look at `rootPath` to see the local directory path on the machine that submitted the job. To see the Amazon S3 object suffix that contains a manifest file, look at `inputManifestFile`. The manifest file contains metadata for a directory tree snapshot of the job's input data.

3. Pretty-print the Amazon S3 manifest object to see the input directory structure for the job.

```
MANIFEST_SUFFIX=$(aws deadline get-job \
    --farm-id $DEV_FARM_ID \
    --queue-id $DEV_QUEUE_ID \
    --job-id $JOB_ID \
    --query "attachments.manifests[0].inputManifestPath" \
    --output text)
  aws s3 cp s3://$DEV_FARM_BUCKET/JobAttachments/Manifests/$MANIFEST_SUFFIX - | jq .
```

Output such as the following is shown:

```
{
    "hashAlg": "xxh128",
    "manifestVersion": "2023-03-03",
    "paths": [
    {
        "hash": "2ec297b04c59c4741ed97ac8fb83080c",
        "mtime": 1698186190000000,
        "path": "simple_job/template.yaml",
        "size": 445
    }
    ],
    "totalSize": 445
  }
```

4. Construct the Amazon S3 prefix that holds manifests for the output job attachments and list the object under it.

```
SESSION_ACTION=$(aws deadline list-session-actions \
    --farm-id $DEV_FARM_ID \
    --queue-id $DEV_QUEUE_ID \
    --job-id $JOB_ID \
    --session-id $SESSION_ID \
    --query "sessionActions[?definition.taskRun != null] | [0]")
STEP_ID=$(echo $SESSION_ACTION | jq -r .definition.taskRun.stepId)
TASK_ID=$(echo $SESSION_ACTION | jq -r .definition.taskRun.taskId)
TASK_OUTPUT_PREFIX=JobAttachments/Manifests/$DEV_FARM_ID/$DEV_QUEUE_ID/$JOB_ID/
$STEP_ID/$TASK_ID/
aws s3api list-objects-v2 --bucket $DEV_FARM_BUCKET --prefix $TASK_OUTPUT_PREFIX
```

The output job attachments are not directly referenced from the job resource but are instead placed in an Amazon S3 bucket based on farm resource IDs.

5. Get the newest manifest object key for the specific session action id, then pretty-print the manifest objects.

```
SESSION_ACTION_ID=$(echo $SESSION_ACTION | jq -r .sessionActionId)
 MANIFEST_KEY=$(aws s3api list-objects-v2 \
    --bucket $DEV_FARM_BUCKET \
    --prefix $TASK_OUTPUT_PREFIX \
    --query "Contents[*].Key" --output text \
    | grep $SESSION_ACTION_ID \
    | sort | tail -1)
 MANIFEST_OBJECT=$(aws s3 cp s3://$DEV_FARM_BUCKET/$MANIFEST_KEY -)
 echo $MANIFEST_OBJECT | jq .
```

You'll see properties of the file hash-jobattachments.txt in the output such as the following:

```
{
    "hashAlg": "xxh128",
    "manifestVersion": "2023-03-03",
    "paths": [
    {
        "hash": "f60b8e7d0fabf7214ba0b6822e82e08b",
        "mtime": 1698785252554950,
        "path": "hash-jobattachments.txt",
        "size": 182
    }
```

```
      ],
      "totalSize": 182
  }
```

Your job will only have a single manifest object per task run, but in general it is possible to have more of objects per task run.

6. View content-addressible Amazon S3 storage output under the `Data` prefix.

```
FILE_HASH=$(echo $MANIFEST_OBJECT | jq -r .paths[0].hash)
FILE_PATH=$(echo $MANIFEST_OBJECT | jq -r .paths[0].path)
aws s3 cp s3://$DEV_FARM_BUCKET/JobAttachments/Data/$FILE_HASH -
```

Output such as the following is shown:

```
eaa2df5d34b54be5ac34c56a24a8c237b8487231a607eaf530a04d76b89c9cd3  /tmp/openjd/
session-123/assetroot-abc/simple_job/template.yaml
```

Next steps

After learning how to submit jobs with attachments using the Deadline Cloud CLI, you can explore:

- Submit with Deadline Cloud to learn how to run jobs using an OpenJD bundle on your worker hosts.

- Add a service-managed fleet to your developer farm in Deadline Cloud to run your jobs on hosts managed by Deadline Cloud.

- Clean up your farm resources in Deadline Cloud to shut down the resources that you used for this tutorial.

Add a service-managed fleet to your developer farm in Deadline Cloud

AWS CloudShell does not provide enough compute capacity to test larger workloads. It's also not configured to work with jobs that distribute tasks on multiple worker hosts.

Instead of using CloudShell,you can add an Auto Scaling service-managed fleet (SMF) to your developer farm. An SMF provides sufficient compute capacity for larger workloads and can handle jobs that need to distribute job tasks across multiple worker hosts. The scheduler will use both the SMF and CMF workers to run jobs, unless you shut down the CMF worker.

Before you add an SMF, you must set up a Deadline Cloud farm, queue, and fleet. See [Create a Deadline Cloud farm](#).

To add a service-managed fleet to your developer farm

1. Choose your first AWS CloudShell tab, then create the service managed fleet and add its fleet ID to `.bashrc`. This action makes it available for other terminal sessions.

```
FLEET_ROLE_ARN="arn:aws:iam::$(aws sts get-caller-identity \
        --query "Account" --output text):role/${DEV_FARM_NAME}FleetRole"
 aws deadline create-fleet \
     --farm-id $DEV_FARM_ID \
     --display-name "$DEV_FARM_NAME SMF" \
     --role-arn $FLEET_ROLE_ARN \
     --max-worker-count 5 \
     --configuration \
        '{
            "serviceManagedEc2": {
                "instanceCapabilities": {
                    "vCpuCount": {
                        "min": 2,
                        "max": 4
                    },
                    "memoryMiB": {
                        "min": 512
                    },
                    "osFamily": "linux",
                    "cpuArchitectureType": "x86_64"
                },
                "instanceMarketOptions": {
                    "type": "spot"
                }
            }
        }'

  echo "DEV_SMF_ID=$(aws deadline list-fleets \
        --farm-id $DEV_FARM_ID \
        --query "fleets[?displayName=='$DEV_FARM_NAME SMF'].fleetId \
```

```
         | [0]" --output text)" >> ~/.bashrc
  source ~/.bashrc
```

2. Associate the SMF with your queue.

```
aws deadline create-queue-fleet-association \
    --farm-id $DEV_FARM_ID \
    --queue-id $DEV_QUEUE_ID \
    --fleet-id $DEV_SMF_ID
```

3.

> **ⓘ Note**
>
> The scheduler will use both the SMF and CMF workers to run jobs, unless you shut
> down the CMF worker.

Submit simple_file_job to the queue. When prompted to confirm the upload, enter **y**.

```
deadline bundle submit simple_file_job \
    -p InFile=simple_job/template.yaml \
    -p OutFile=hash-jobattachments.txt
```

4. Confirm the SMF is working correctly.

```
deadline fleet get
```

- The worker may take a few minutes to start.

- The queueFleetAssociationsStatus for your customer managed fleet and service
  managed fleet will be ACTIVE.

- The SMF autoScalingStatus will change from GROWING to STEADY.

Your status will look similar to the following:

```
fleetId: fleet-2cc78e0dd3f04d1db427e7dc1d51ea44
farmId: farm-63ee8d77cdab4a578b685be8c5561c4a
displayName: DeveloperFarm SMF
description: ''
status: ACTIVE
autoScalingStatus: STEADY
```

```
    targetWorkerCount: 0
    workerCount: 0
    minWorkerCount: 0
    maxWorkerCount: 5
```

5. View the log for the job that you submitted. This log is stored in a log in Amazon CloudWatch Logs, not the CloudShell file system.

```
    JOB_ID=$(deadline config get defaults.job_id)
    SESSION_ID=$(aws deadline list-sessions \
          --farm-id $DEV_FARM_ID \
          --queue-id $DEV_QUEUE_ID \
          --job-id $JOB_ID \
          --query "sessions[0].sessionId" \
          --output text)
    aws logs tail /aws/deadline/$DEV_FARM_ID/$DEV_QUEUE_ID \
        --log-stream-names $SESSION_ID
```

Next steps

After creating and testing a service-managed fleet, you should remove the resources that you created to avoid unnecessary charges.

*   [Clean up your farm resources in Deadline Cloud](#) to shut down the resources that you used for this tutorial.

Clean up your farm resources in Deadline Cloud

To develop and test new workloads and pipeline integrations, you can continue to use the Deadline Cloud developer farm that you created for this tutorial. If you no longer need your developer farm, you can delete its resources including farm, fleet, queue, AWS Identity and Access Management (IAM) roles, and logs in Amazon CloudWatch Logs. After you delete these resources, you will need to begin the tutorial again to use the resources. For more information, see [Getting started with Deadline Cloud resources.](#).

To clean up developer farm resources

1. Choose your first CloudShell tab, then stop all the queue-fleet associations for your queue.

```
FLEETS=$(aws deadline list-queue-fleet-associations \
        --farm-id $DEV_FARM_ID \
        --queue-id $DEV_QUEUE_ID \
        --query "queueFleetAssociations[].fleetId" \
        --output text)
for FLEET_ID in $FLEETS; do
    aws deadline update-queue-fleet-association \
        --farm-id $DEV_FARM_ID \
        --queue-id $DEV_QUEUE_ID \
        --fleet-id $FLEET_ID \
        --status STOP_SCHEDULING_AND_CANCEL_TASKS
done
```

2. List the queue fleet associations.

```
aws deadline list-queue-fleet-associations \
      --farm-id $DEV_FARM_ID \
      --queue-id $DEV_QUEUE_ID
```

You might need to rerun the command until the output reports "status": "STOPPED", then you can proceed to the next step. This process can take several minutes to complete.

```
{
    "queueFleetAssociations": [
        {
            "queueId": "queue-abcdefgh01234567890123456789012id",
            "fleetId": "fleet-abcdefgh01234567890123456789012id",
            "status": "STOPPED",
            "createdAt": "2023-11-21T20:49:19+00:00",
            "createdBy": "arn:aws:sts::123456789012:assumed-role/RoleToBeAssumed/
MySessionName",
            "updatedAt": "2023-11-21T20:49:38+00:00",
            "updatedBy": "arn:aws:sts::123456789012:assumed-role/RoleToBeAssumed/
MySessionName"
        },
        {
            "queueId": "queue-abcdefgh01234567890123456789012id",
            "fleetId": "fleet-abcdefgh01234567890123456789012id",
            "status": "STOPPED",
            "createdAt": "2023-11-21T20:32:06+00:00",
```

```
                "createdBy": "arn:aws:sts::123456789012:assumed-role/RoleToBeAssumed/
    MySessionName",
                "updatedAt": "2023-11-21T20:49:39+00:00",
                "updatedBy": "arn:aws:sts::123456789012:assumed-role/RoleToBeAssumed/
    MySessionName"
            }
        ]
    }
```

3. Delete all of the queue-fleet associations for your queue.

```
for FLEET_ID in $FLEETS; do
    aws deadline delete-queue-fleet-association \
        --farm-id $DEV_FARM_ID \
        --queue-id $DEV_QUEUE_ID \
        --fleet-id $FLEET_ID
  done
```

4. Delete all of the fleets associated with your queue.

```
for FLEET_ID in $FLEETS; do
    aws deadline delete-fleet \
        --farm-id $DEV_FARM_ID \
        --fleet-id $FLEET_ID
  done
```

5. Delete the queue.

```
aws deadline delete-queue \
    --farm-id $DEV_FARM_ID \
    --queue-id $DEV_QUEUE_ID
```

6. Delete the farm.

```
aws deadline delete-farm \
    --farm-id $DEV_FARM_ID
```

7. Delete other AWS resources for your farm.

    a. Delete the fleet AWS Identity and Access Management (IAM) role.

    ```
    aws iam delete-role-policy \
        --role-name "${DEV_FARM_NAME}FleetRole" \
    ```

```
        --policy-name WorkerPermissions
aws iam delete-role \
        --role-name "${DEV_FARM_NAME}FleetRole"
```

b. Delete the queue IAM role.

```
aws iam delete-role-policy \
        --role-name "${DEV_FARM_NAME}QueueRole" \
        --policy-name S3BucketsAccess
aws iam delete-role \
        --role-name "${DEV_FARM_NAME}QueueRole"
```

c. Delete the Amazon CloudWatch Logs log groups. Each queue and fleet has their own log
    group.

```
aws logs delete-log-group \
        --log-group-name "/aws/deadline/$DEV_FARM_ID/$DEV_QUEUE_ID"
aws logs delete-log-group \
        --log-group-name "/aws/deadline/$DEV_FARM_ID/$DEV_CMF_ID"
aws logs delete-log-group \
        --log-group-name "/aws/deadline/$DEV_FARM_ID/$DEV_SMF_ID"
```

How to submit a job to Deadline Cloud

There are many different ways to submit jobs to AWS Deadline Cloud. This section describes some of the ways that you can submit jobs using the tools provided by Deadline Cloud or by creating your own custom tools for your workloads.

- From a terminal – for when you're first developing a job bundle, or when users submitting a job are comfortable using the command line
- From a script – for customizing and automating workloads
- From an application – for when the user's work is in an application, or when an application's context is important.

The following examples use the `deadline` Python library and the `deadline` command line tool. Both are available from [PyPi](#) and [hosted on GitHub](#).

Topics

- [Submit a job to Deadline Cloud from a terminal](#)
- [Submit a job to Deadline Cloud using a script](#)
- [Submit a job within an application](#)

Submit a job to Deadline Cloud from a terminal

Using only a job bundle and the Deadline Cloud CLI, you or your more technical users can rapidly iterate on writing job bundles to test submitting a job. Use the following command to submit a job bundle:

```
deadline bundle submit <path-to-job-bundle>
```

If you submit a job bundle with parameters that do not have defaults in the bundle, you can specify them with the `-p` / `--parameter` option.

```
deadline bundle submit <path-to-job-bundle> -p <parameter-name>=<parameter-value> -
p ...
```

For a complete list of the available options, run the help command:

```
deadline bundle submit --help
```

Submit a job to Deadline Cloud using a GUI

The Deadline Cloud CLI also comes with a graphical user interface that enables users to see the parameters they must provide before submitting a job. If your users prefer not to interact with the command line, you can write a desktop shortcut that opens a dialog to submit a specific job bundle:

```
deadline bundle gui-submit <path-to-job-bundle>
```

Use the `--browse` option can so the user can select a job bundle:

```
deadline bundle gui-submit --browse
```

For a complete list of available options, run the help command:

```
deadline bundle gui-submit --help
```

Submit a job to Deadline Cloud using a script

To automate submitting jobs to Deadline Cloud, you can script them using tools such as bash, Powershell, and batch files.

You can add functionality like populating job parameters from environment variables or other applications. You can also submit multiple jobs in a row, or script the creation of a job bundle to submit.

Submit a job using Python

Deadline Cloud also an open-source Python library to interact with the service. The source code is available on GitHub.

The library is available on pypi via pip (`pip install deadline`). It's the same library used by the Deadline Cloud CLI tool:

```
from deadline.client import api
```

```
job_bundle_path = "/path/to/job/bundle"
job_parameters = [
    {
        "name": "parameter_name",
        "value": "parameter_value"
    },
]

job_id = api.create_job_from_job_bundle(
    job_bundle_path,
    job_parameters
)
print(job_id)
```

To create a dialog like the `deadline bundle gui-submit` command, you can use of `show_job_bundle_submitter` function from the `deadline.client.ui.job_bundle_submitter.`

The following example starts a Qt application and shows the job bundle submitter:

```
# The GUI components must be installed with pip install "deadline[gui]"
import sys
from qtpy.QtWidgets import QApplication
from deadline.client.ui.job_bundle_submitter import show_job_bundle_submitter

app = QApplication(sys.argv)
submitter = show_job_bundle_submitter(browse=True)
submitter.show()
app.exec()
print(submitter.create_job_response)
```

To make your own dialog you can use the `SubmitJobToDeadlineDialog` class in `deadline.client.ui.dialogs.submit_job_to_deadline_dialog`. You can pass in values, embed your own job specific tab, and determine how the job bundle gets created (or passed in).

Submit a job within an application

To make it easy for users to submit jobs, you can use the scripting runtimes or plugin systems provided by an application. Users have a familiar interface and you can create powerful tools that assist the users when submitting a workload.

Embed job bundles in an application

This example demonstrates submitting job bundles that you make available in the application.

To give a user access to these job bundles, create a script embedded in a menu item that launches the Deadline Cloud CLI.

The following script enables a user to select the job bundle:

```
deadline bundle gui-submit --install-gui
```

To use a specific job bundle in a menu item instead, use the following:

```
deadline bundle gui-submit </path/to/job/bundle> --install-gui
```

This opens a dialog where the user can modify the job parameters, inputs, and outputs, and then submit the job. You can have different menu items for different job bundles for a user to submit in an application.

If the job that you submit with a job bundle contains similar parameters and asset references across submissions, you can fill in the default values in the underlying job bundle.

Get information from an application

To pull information from an application so that users don't have to manually add it to the submission, you can integrate Deadline Cloud with the application so that your users can submit jobs using a familiar interface without needing exit the application or use command line tools.

If your application has a scripting runtime that supports Python and pyside/pyqt, you can use the GUI components from the [Deadline Cloud client library](#) to create a UI. For an example, see [Deadline Cloud for Maya integration](#) on GitHub.

The Deadline Cloud client library provides operations that do the following to help you provide a strong integrated user experience:

- Pull queue environment parameters, job parameters, and asset references form environment variables and by calling the application SDK.
- Set the parameters in the job bundle. To avoid modifying the original bundle, you should make a copy of the bundle and submit the copy.

If you use the `deadline bundle gui-submit` command to submit the job bundle, you must programmatically the `parameter_values.yaml` and `asset_references.yaml` files to pass the information from the application. For more information about these files see Open Job Description (OpenJD) templates for Deadline Cloud.

If you need more complex controls than the ones offered by OpenJD, need to abstract the job from the user, or want to make the integration match the application's visual style, you can write your own dialog that calls the Deadline Cloud client library to submit the job.

Configure jobs using queue environments

AWS Deadline Cloud uses *queue environments* to configure the software on your workers. An environment enables you to perform time-consuming tasks, such as set up and tear-down, once for all the tasks in a session. It defines the actions to run on a worker when starting or stopping a session. You can configure an environment for a queue, jobs that run in the queue, and the individual steps for a job.

You define environments as queue environments or job environments. Create queue environments with the Deadline Cloud console or with the [deadline:CreateQueueEnvironment](#) operation and define job environments in the job templates of the jobs you submit. They follow the Open Job Description (OpenJD) specification for environments. For details, see [<Environment>](#) in the OpenJD specification on GitHub.

In addition to a `name` and `description`, each environment contains two fields that define the environment on the host. They are:

- `script` – The action taken when this environment is run on a worker.

- `variables` – A set of environment variable name/value pairs that are set when entering the environment.

You must set at least one of `script` or `variables`.

You can define more than one environment in your job template. Each environment is applied in the order that they are listed in the template. You can use this to help manage the complexity of your environments.

The default queue environment for Deadline Cloud uses the conda package manager to load software into the environment, but you can use other package managers. The default environment defines two parameters to specify the software that should be loaded. These variables are set by submitters provided by Deadline Cloud, though you can set them in your own scripts and applications that use the default environment. They are:

- `CondaPackages` – A space-separated list of conda package match specifications to install for the job. For example, the Blender submitter would add `blender=3.6` to render frames in Blender 3.6.

- `CondaChannels` – A space-separated list of conda channels to install packages from. For service-managed fleets, packages are installed from the `deadline-cloud` channel. You can add other channels.

Topics

- [Control the job environment with OpenJD queue environments](#)

- [Provide applications for your jobs](#)

- [Create a conda channel using S3](#)

- [Create a conda package for an application](#)

Control the job environment with OpenJD queue environments

You can define customized environments for your rendering jobs using *queue environments*. A queue environment is a template that controls the environment variables, file mappings, and other settings for jobs running in a specific queue. It enables you to tailor the execution environment for the jobs submitted to a queue to the requirements of your workloads. AWS Deadline Cloud provides three nested levels where you can apply [Open Job Description (OpenJD) environments](#): queue, job, and step. By defining queue environments, you can ensure consistent and optimized performance for different types of jobs, streamline resource allocation, and simplify queue management.

The queue environment is a template that you attach to a queue in your AWS account from the AWS management console or using the AWS CLI. You can create one environment for a queue, or you can create multiple queue environments that applied in order to create the execution environment. This enables you to create and test an environment in steps to help ensure that it works correctly for you jobs.

Job and step environments are defined in the job template you use to create a job in your queue. The OpenJD syntax is the same in these different forms of environments. In this section we will show them inside of job templates.

Topics

- [Set environment variables in a queue environment](#)

- [Set the path in a queue environment](#)

- [Run a background daemon process from the queue environment](#)

Set environment variables in a queue environment

[Open Job Description (OpenJD) environments](#) can set environment variables that every task command within their scope uses. Many applications and frameworks check for environment variables to control feature settings, logging level, and more.

For example, the [Qt Framework](#) provides GUI functionality for many desktop applications. When you run these applications on a worker host without an interactive display, you may need to set the environment variable QT_QPA_PLATFORM to offscreen so the worker doesn't look for a display.

In this example, you'll use a sample job bundle from the Deadline Cloud samples directory to set and view the environment variables for a job.

Prerequisites

Perform the following steps to run the [sample job bundle with environment variables](#) from the Deadline Cloud samples github repository.

1. If you do not have a Deadline Cloud farm with a queue and associated Linux fleet, follow the guided onboarding experience in the [Deadline Cloud console](#) to create one with default settings.

2. If you do not have the Deadline Cloud CLI and Deadline Cloud monitor on your workstation, follow the steps in [Set up Deadline Cloud submitters](#) from the user guide.

3. Use git to clone the [Deadline Cloud samples GitHub repository](#).

```
git clone https://github.com/aws-deadline/deadline-cloud-samples.git
  Cloning into 'deadline-cloud-samples'...
  ...
cd deadline-cloud-samples/job_bundles
```

Run the environment variable sample

1. Use the Deadline Cloud CLI to submit the job_env_vars sample.

```
deadline bundle submit job_env_vars
  Submitting to Queue: MySampleQueue
  ...
```

2. In the Deadline Cloud monitor, you can see the new job and monitor its progress. After the Linux fleet associated with the queue has a worker available to run the job's task, the job completes

in a few seconds. Select the task, then choose the **View logs** option in the top right menu of the tasks panel.

On the right are three session actions, **Launch JobEnv**, **Launch StepEnv**, and **Task run**. The log view in the center of the window corresponds to the selected session action on the right.

Compare the session actions with their definitions

In this section you use the Deadline Cloud monitor to compare the session actions with where they are defined in the job template. It continues from the previous section.

Open the file job_env_vars/template.yaml in a text editor. This is the job template that defines the session actions.

1. Select the **Launch JobEnv** session action in Deadline Cloud monitor. You will see the following log output.

```
024/07/16 16:18:27-07:00
2024/07/16 16:18:27-07:00 ===========================================
2024/07/16 16:18:27-07:00 --------- Entering Environment: JobEnv
2024/07/16 16:18:27-07:00 ===========================================
2024/07/16 16:18:27-07:00 Setting: JOB_VERBOSITY=MEDIUM
2024/07/16 16:18:27-07:00 Setting: JOB_EXAMPLE_PARAM=An example parameter value
2024/07/16 16:18:27-07:00 Setting: JOB_PROJECT_ID=project-12
2024/07/16 16:18:27-07:00 Setting: JOB_ENDPOINT_URL=https://internal-host-name/some/
path
2024/07/16 16:18:27-07:00 Setting: QT_QPA_PLATFORM=offscreen
```

The following lines from the job template specified this action.

```
jobEnvironments:
 - name: JobEnv
   description: Job environments apply to everything in the job.
   variables:
     # When applications have options as environment variables, you can set them
 here.
     JOB_VERBOSITY: MEDIUM
     # You can use the value of job parameters when setting environment variables.
     JOB_EXAMPLE_PARAM: "{{Param.ExampleParam}}"
     # Some more ideas.
     JOB_PROJECT_ID: project-12
```

```
JOB_ENDPOINT_URL: https://internal-host-name/some/path
# This variable lets applications using the Qt Framework run without a display
QT_QPA_PLATFORM: offscreen
```

2. Select the **Launch StepEnv** session action in Deadline Cloud monitor. You will see the following log output.

```
2024/07/16 16:18:27-07:00
2024/07/16 16:18:27-07:00 ==========================================
2024/07/16 16:18:27-07:00 --------- Entering Environment: StepEnv
2024/07/16 16:18:27-07:00 ==========================================
2024/07/16 16:18:27-07:00 Setting: STEP_VERBOSITY=HIGH
2024/07/16 16:18:27-07:00 Setting: JOB_PROJECT_ID=step-project-12
```

The following lines from the job template specified this action.

```
stepEnvironments:
- name: StepEnv
  description: Step environments apply to all the tasks in the step.
  variables:
    # These environment variables are only set within this step, not other steps.
    STEP_VERBOSITY: HIGH
    # Replace a variable value defined at the job level.
    JOB_PROJECT_ID: step-project-12
```

3. Select the **Task run** session action in Deadline Cloud monitor. You will see the following output.

```
2024/07/16 16:18:27-07:00
2024/07/16 16:18:27-07:00 ==========================================
2024/07/16 16:18:27-07:00 --------- Running Task
2024/07/16 16:18:27-07:00 ==========================================
2024/07/16 16:18:27-07:00 ------------------------------------------------
2024/07/16 16:18:27-07:00 Phase: Setup
2024/07/16 16:18:27-07:00 ------------------------------------------------
2024/07/16 16:18:27-07:00 Writing embedded files for Task to disk.
2024/07/16 16:18:27-07:00 Mapping: Task.File.Run -> /sessions/session-
b4bd451784674c0987be82c5f7d5642deupf6tk9/embedded_files08cdnuyt/tmpmdiajwvh
2024/07/16 16:18:27-07:00 Wrote: Run -> /sessions/session-
b4bd451784674c0987be82c5f7d5642deupf6tk9/embedded_files08cdnuyt/tmpmdiajwvh
2024/07/16 16:18:27-07:00 ------------------------------------------------
2024/07/16 16:18:27-07:00 Phase: Running action
2024/07/16 16:18:27-07:00 ------------------------------------------------
```

```
  2024/07/16 16:18:27-07:00 Running command sudo -u job-user -i setsid -w /sessions/
 session-b4bd451784674c0987be82c5f7d5642deupf6tk9/tmpiqbrsby4.sh
 2024/07/16 16:18:27-07:00 Command started as pid: 2176
 2024/07/16 16:18:27-07:00 Output:
 2024/07/16 16:18:28-07:00 Running the task
 2024/07/16 16:18:28-07:00
 2024/07/16 16:18:28-07:00 Environment variables starting with JOB_*:
 2024/07/16 16:18:28-07:00 JOB_ENDPOINT_URL=https://internal-host-name/some/path
 2024/07/16 16:18:28-07:00 JOB_EXAMPLE_PARAM='An example parameter value'
 2024/07/16 16:18:28-07:00 JOB_PROJECT_ID=step-project-12
 2024/07/16 16:18:28-07:00 JOB_VERBOSITY=MEDIUM
 2024/07/16 16:18:28-07:00
 2024/07/16 16:18:28-07:00 Environment variables starting with STEP_*:
 2024/07/16 16:18:28-07:00 STEP_VERBOSITY=HIGH
 2024/07/16 16:18:28-07:00
 2024/07/16 16:18:28-07:00 Done running the task
 2024/07/16 16:18:28-07:00 ----------------------------------------------
 2024/07/16 16:18:28-07:00 Uploading output files to Job Attachments
 2024/07/16 16:18:28-07:00 ----------------------------------------------
```

The following lines from the job template specified this action.

```
  script:
    actions:
      onRun:
        command: bash
        args:
        - '{{Task.File.Run}}'
    embeddedFiles:
    - name: Run
      type: TEXT
      data: |
        echo Running the task
        echo ""

        echo Environment variables starting with JOB_*:
        set | grep ^JOB_
        echo ""

        echo Environment variables starting with STEP_*:
        set | grep ^STEP_
        echo ""
```

```
                echo Done running the task
```

Set the path in a queue environment

Use OpenJD environments to provide new commands in an environment. First you create a directory containing script files, and then add that directory to the PATH environment variables so that executables in your script can run them without needing to specify the directory path each time. The list of variables in an environment definition doesn't provide a way to modify the variable, so you do this by running a script instead. After the script sets things up and modifies the PATH, it exports the variable to the OpenJD runtime with the command `echo "openjd_env: PATH=$PATH"`.

Prerequisites

Perform the following steps to run the [sample job bundle with environment variables](#) from the Deadline Cloud samples github repository.

1. If you do not have a Deadline Cloud farm with a queue and associated Linux fleet, follow the guided onboarding experience in the [Deadline Cloud console](#) to create one with default settings.
2. If you do not have the Deadline Cloud CLI and Deadline Cloud monitor on your workstation, follow the steps in [Set up Deadline Cloud submitters](#) from the user guide.
3. Use `git` to clone the [Deadline Cloud samples GitHub repository](#).

   ```
   git clone https://github.com/aws-deadline/deadline-cloud-samples.git
     Cloning into 'deadline-cloud-samples'...
     ...
   cd deadline-cloud-samples/job_bundles
   ```

Run the path sample

1. Use the Deadline Cloud CLI to submit the `job_env_with_new_command` sample.

   ```
   $ deadline bundle submit job_env_with_new_command
   Submitting to Queue: MySampleQueue
   ...
   ```

2. In the Deadline Cloud monitor, you will see the new job and can monitor its progress. Once the Linux fleet associated with the queue has a worker available to run the job's task, the job

completes in a few seconds. Select the task, then choose the **View logs** option in the top right menu of the tasks panel.

On the right are two session actions, **Launch RandomSleepCommand** and **Task run**. The log viewer in the center of the window corresponds to the selected session action on the right.

Compare session actions with their definitions

In this section you use the Deadline Cloud monitor to compare the session actions with where they are defined in the job template. It continues from the previous section.

Open the file job_env_with_new_command/template.yaml in a text editor. Compare the session actions to where they are defined in the job template.

1. Select the **Launch RandomSleepCommand** session action in the Deadline Cloud monitor. You will see log output as follows.

```
2024/07/16 17:25:32-07:00
2024/07/16 17:25:32-07:00 ==============================================
2024/07/16 17:25:32-07:00 --------- Entering Environment: RandomSleepCommand
2024/07/16 17:25:32-07:00 ==============================================
2024/07/16 17:25:32-07:00 ---------------------------------------------
2024/07/16 17:25:32-07:00 Phase: Setup
2024/07/16 17:25:32-07:00 ---------------------------------------------
2024/07/16 17:25:32-07:00 Writing embedded files for Environment to disk.
2024/07/16 17:25:32-07:00 Mapping: Env.File.Enter -> /sessions/session-
ab132a51b9b54d5da22cbe839dd946baaw1c8hk5/embedded_filesf3tq_1os/tmpbt8j_c3f
2024/07/16 17:25:32-07:00 Mapping: Env.File.SleepScript -> /sessions/session-
ab132a51b9b54d5da22cbe839dd946baaw1c8hk5/embedded_filesf3tq_1os/tmperastlp4
2024/07/16 17:25:32-07:00 Wrote: Enter -> /sessions/session-
ab132a51b9b54d5da22cbe839dd946baaw1c8hk5/embedded_filesf3tq_1os/tmpbt8j_c3f
2024/07/16 17:25:32-07:00 Wrote: SleepScript -> /sessions/session-
ab132a51b9b54d5da22cbe839dd946baaw1c8hk5/embedded_filesf3tq_1os/tmperastlp4
2024/07/16 17:25:32-07:00 ---------------------------------------------
2024/07/16 17:25:32-07:00 Phase: Running action
2024/07/16 17:25:32-07:00 ---------------------------------------------
2024/07/16 17:25:32-07:00 Running command sudo -u job-user -i setsid -w /sessions/
session-ab132a51b9b54d5da22cbe839dd946baaw1c8hk5/tmpbwrquq5u.sh
2024/07/16 17:25:32-07:00 Command started as pid: 2205
2024/07/16 17:25:32-07:00 Output:
2024/07/16 17:25:33-07:00 openjd_env: PATH=/sessions/session-
ab132a51b9b54d5da22cbe839dd946baaw1c8hk5/bin:/opt/conda/condabin:/home/job-
```

```
user/.local/bin:/home/job-user/bin:/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/
bin:/sbin:/bin:/var/lib/snapd/snap/bin
 No newer logs at this moment.
```

The following lines from the job template specified this action.

```
jobEnvironments:
- name: RandomSleepCommand
  description: Adds a command 'random-sleep' to the environment.
  script:
    actions:
      onEnter:
        command: bash
        args:
        - "{{Env.File.Enter}}"
    embeddedFiles:
    - name: Enter
      type: TEXT
      data: |
        #!/bin/env bash
        set -euo pipefail

        # Make a bin directory inside the session's working directory for providing
new commands
        mkdir -p '{{Session.WorkingDirectory}}/bin'

        # If this bin directory is not already in the PATH, then add it
        if ! [[ ":$PATH:" == *':{{Session.WorkingDirectory}}/bin:'* ]]; then
          export "PATH={{Session.WorkingDirectory}}/bin:$PATH"

          # This message to Open Job Description exports the new PATH value to the
environment
          echo "openjd_env: PATH=$PATH"
        fi

        # Copy the SleepScript embedded file into the bin directory
        cp '{{Env.File.SleepScript}}' '{{Session.WorkingDirectory}}/bin/random-
sleep'
        chmod u+x '{{Session.WorkingDirectory}}/bin/random-sleep'
    - name: SleepScript
      type: TEXT
      runnable: true
      data: |
```

```
...
```

2. Select the **Launch StepEnv** session action in the Deadline Cloud monitor. You see log output as follows.

```
2024/07/16 17:25:33-07:00
2024/07/16 17:25:33-07:00 ==========================================
2024/07/16 17:25:33-07:00 --------- Running Task
2024/07/16 17:25:33-07:00 ==========================================
2024/07/16 17:25:33-07:00 ------------------------------------------
2024/07/16 17:25:33-07:00 Phase: Setup
2024/07/16 17:25:33-07:00 ------------------------------------------
2024/07/16 17:25:33-07:00 Writing embedded files for Task to disk.
2024/07/16 17:25:33-07:00 Mapping: Task.File.Run -> /sessions/session-
ab132a51b9b54d5da22cbe839dd946baaw1c8hk5/embedded_filesf3tq_1os/tmpdrwuehjf
2024/07/16 17:25:33-07:00 Wrote: Run -> /sessions/session-
ab132a51b9b54d5da22cbe839dd946baaw1c8hk5/embedded_filesf3tq_1os/tmpdrwuehjf
2024/07/16 17:25:33-07:00 ------------------------------------------
2024/07/16 17:25:33-07:00 Phase: Running action
2024/07/16 17:25:33-07:00 ------------------------------------------
2024/07/16 17:25:33-07:00 Running command sudo -u job-user -i setsid -w /sessions/
session-ab132a51b9b54d5da22cbe839dd946baaw1c8hk5/tmpz81iaqfw.sh
2024/07/16 17:25:33-07:00 Command started as pid: 2256
2024/07/16 17:25:33-07:00 Output:
2024/07/16 17:25:34-07:00 + random-sleep 12.5 27.5
2024/07/16 17:26:00-07:00 Sleeping for duration 26.90
2024/07/16 17:26:00-07:00 ------------------------------------------
2024/07/16 17:26:00-07:00 Uploading output files to Job Attachments
2024/07/16 17:26:00-07:00 ------------------------------------------
```

3. The following lines from the job template specified this action.

```
steps:
- name: EnvWithCommand
  script:
    actions:
      onRun:
        command: bash
        args:
        - '{{Task.File.Run}}'
    embeddedFiles:
    - name: Run
      type: TEXT
```

```
        data: |
          set -xeuo pipefail

          # Run the script installed into PATH by the job environment
          random-sleep 12.5 27.5
    hostRequirements:
      attributes:
      - name: attr.worker.os.family
        anyOf:
        - linux
```

Run a background daemon process from the queue environment

In many rendering use cases, loading the application and scene data can take a significant amount of time. If a job reloads them for every frame, it will spend most of its time on overhead. It's often possible to load the application once as a background daemon process, have it load the scene data, and then send it commands via inter-process communication (IPC) to perform the renders.

Many of the open source Deadline Cloud integrations use this pattern. The Open Job Description project provides an [adaptor runtime library](#) with robust IPC patterns on all supported operating systems.

To demonstrate this pattern, there is a [self-contained sample job bundle](#) that uses Python and bash code to implement a background daemon and the IPC for tasks to communicate with it. The daemon is implemented in Python, and listens for a POSIX SIGUSR1 signal for when to process a task. The task details are passed to the daemon in a specific JSON file, and the results of running the task are returned as another JSON file.

Prerequisites

Perform the following steps to run the [sample job bundle with a daemon process](#) from the Deadline Cloud samples github repository.

1. If you do not have a Deadline Cloud farm with a queue and associated Linux fleet, follow the guided onboarding experience in the [Deadline Cloud console](#) to create one with default settings.
2. If you do not have the Deadline Cloud CLI and Deadline Cloud monitor on your workstation, follow the steps in [Set up Deadline Cloud submitters](#) from the user guide.
3. Use `git` to clone the [Deadline Cloud samples GitHub repository](#).

```
git clone https://github.com/aws-deadline/deadline-cloud-samples.git
```

```
  Cloning into 'deadline-cloud-samples'...
  ...
cd deadline-cloud-samples/job_bundles
```

Run the daemon sample

1. Use the Deadline Cloud CLI to submit the `job_env_daemon_process` sample.

```
  git clone https://github.com/aws-deadline/deadline-cloud-samples.git
Cloning into 'deadline-cloud-samples'...
  ...
cd deadline-cloud-samples/job_bundles
```

2. In the Deadline Cloud monitor application, you will see the new job and can monitor its progress. Once the Linux fleet associated with the queue has a worker available to run the job's task, it completes in about a minute. With one of the tasks selected, choose the **View logs** option in the top right menu of the tasks panel.

   On the right there are two session actions, **Launch DaemonProcess** and **Task run**. The log viewer in the center of the window corresponds to the selected session action on the right.

   Select the option **View logs for all tasks**. The timeline shows the rest of the tasks that ran as part of the session, and the `Shut down DaemonProcess` action that exited the environment.

View the daemon logs

1. In this section you use the Deadline Cloud monitor to compare the session actions with where they are defined in the job template. It continues from the previous section.

   Open the file job_env_daemon_process/template.yaml in a text editor. Compare the session actions to where they are defined in the job template.

2. Select the `Launch DaemonProcess` session action in Deadline Cloud monitor. You will see log output as follows.

```
  2024/07/17 16:27:20-07:00
  2024/07/17 16:27:20-07:00 ==========================================
  2024/07/17 16:27:20-07:00 --------- Entering Environment: DaemonProcess
  2024/07/17 16:27:20-07:00 ==========================================
  2024/07/17 16:27:20-07:00 ------------------------------------------
```

```
 2024/07/17 16:27:20-07:00 Phase: Setup
 2024/07/17 16:27:20-07:00 ---------------------------------------------
 2024/07/17 16:27:20-07:00 Writing embedded files for Environment to disk.
 2024/07/17 16:27:20-07:00 Mapping: Env.File.Enter -> /sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/embedded_fileswy00x5ra/enter-daemon-
process-env.sh
 2024/07/17 16:27:20-07:00 Mapping: Env.File.Exit -> /sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/embedded_fileswy00x5ra/exit-daemon-
process-env.sh
 2024/07/17 16:27:20-07:00 Mapping: Env.File.DaemonScript -> /sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/embedded_fileswy00x5ra/daemon-
script.py
 2024/07/17 16:27:20-07:00 Mapping: Env.File.DaemonHelperFunctions -> /sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/embedded_fileswy00x5ra/daemon-
helper-functions.sh
 2024/07/17 16:27:20-07:00 Wrote: Enter -> /sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/embedded_fileswy00x5ra/enter-daemon-
process-env.sh
 2024/07/17 16:27:20-07:00 Wrote: Exit -> /sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/embedded_fileswy00x5ra/exit-daemon-
process-env.sh
 2024/07/17 16:27:20-07:00 Wrote: DaemonScript -> /sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/embedded_fileswy00x5ra/daemon-
script.py
 2024/07/17 16:27:20-07:00 Wrote: DaemonHelperFunctions -> /sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/embedded_fileswy00x5ra/daemon-
helper-functions.sh
 2024/07/17 16:27:20-07:00 ---------------------------------------------
 2024/07/17 16:27:20-07:00 Phase: Running action
 2024/07/17 16:27:20-07:00 ---------------------------------------------
 2024/07/17 16:27:20-07:00 Running command sudo -u job-user -i setsid -w /sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/tmp_u8slys3.sh
 2024/07/17 16:27:20-07:00 Command started as pid: 2187
 2024/07/17 16:27:20-07:00 Output:
 2024/07/17 16:27:21-07:00 openjd_env: DAEMON_LOG=/sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/daemon.log
 2024/07/17 16:27:21-07:00 openjd_env: DAEMON_PID=2223
 2024/07/17 16:27:21-07:00 openjd_env: DAEMON_BASH_HELPER_SCRIPT=/sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/embedded_fileswy00x5ra/daemon-
helper-functions.sh
```

The following lines from the job template specified this action.

```
    stepEnvironments:
    - name: DaemonProcess
      description: Runs a daemon process for the step's tasks to share.
      script:
        actions:
          onEnter:
            command: bash
            args:
            - "{{Env.File.Enter}}"
          onExit:
            command: bash
            args:
            - "{{Env.File.Exit}}"
        embeddedFiles:
        - name: Enter
          filename: enter-daemon-process-env.sh
          type: TEXT
          data: |
            #!/bin/env bash
            set -euo pipefail

            DAEMON_LOG='{{Session.WorkingDirectory}}/daemon.log'
            echo "openjd_env: DAEMON_LOG=$DAEMON_LOG"
            nohup python {{Env.File.DaemonScript}} > $DAEMON_LOG 2>&1 &
            echo "openjd_env: DAEMON_PID=$!"
            echo "openjd_env:
  DAEMON_BASH_HELPER_SCRIPT={{Env.File.DaemonHelperFunctions}}"

            echo 0 > 'daemon_log_cursor.txt'
      ...
```

3. Select one of the Task run: N session action in Deadline Cloud monitor. You will see log output as follows.

```
2024/07/17 16:27:22-07:00
 2024/07/17 16:27:22-07:00 ===========================================
 2024/07/17 16:27:22-07:00 --------- Running Task
 2024/07/17 16:27:22-07:00 ===========================================
 2024/07/17 16:27:22-07:00 Parameter values:
 2024/07/17 16:27:22-07:00 Frame(INT) = 2
 2024/07/17 16:27:22-07:00 -------------------------------------------
 2024/07/17 16:27:22-07:00 Phase: Setup
```

```
 2024/07/17 16:27:22-07:00 ---------------------------------------------
 2024/07/17 16:27:22-07:00 Writing embedded files for Task to disk.
 2024/07/17 16:27:22-07:00 Mapping: Task.File.Run -> /sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/embedded_fileswy00x5ra/run-task.sh
 2024/07/17 16:27:22-07:00 Wrote: Run -> /sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/embedded_fileswy00x5ra/run-task.sh
 2024/07/17 16:27:22-07:00 ---------------------------------------------
 2024/07/17 16:27:22-07:00 Phase: Running action
 2024/07/17 16:27:22-07:00 ---------------------------------------------
 2024/07/17 16:27:22-07:00 Running command sudo -u job-user -i setsid -w /sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/tmpv4obfkhn.sh
 2024/07/17 16:27:22-07:00 Command started as pid: 2301
 2024/07/17 16:27:22-07:00 Output:
 2024/07/17 16:27:23-07:00 Daemon PID is 2223
 2024/07/17 16:27:23-07:00 Daemon log file is /sessions/
session-972e21d98dde45e59c7153bd9258a64dohwg4yg1/daemon.log
 2024/07/17 16:27:23-07:00
 2024/07/17 16:27:23-07:00 === Previous output from daemon
 2024/07/17 16:27:23-07:00 ===
 2024/07/17 16:27:23-07:00
 2024/07/17 16:27:23-07:00 Sending command to daemon
 2024/07/17 16:27:23-07:00 Received task result:
 2024/07/17 16:27:23-07:00 {
 2024/07/17 16:27:23-07:00   "result": "SUCCESS",
 2024/07/17 16:27:23-07:00   "processedTaskCount": 1,
 2024/07/17 16:27:23-07:00   "randomValue": 0.2578537967668988,
 2024/07/17 16:27:23-07:00   "failureRate": 0.1
 2024/07/17 16:27:23-07:00 }
 2024/07/17 16:27:23-07:00
 2024/07/17 16:27:23-07:00 === Daemon log from running the task
 2024/07/17 16:27:23-07:00 Loading the task details file
 2024/07/17 16:27:23-07:00 Received task details:
 2024/07/17 16:27:23-07:00 {
 2024/07/17 16:27:23-07:00  "pid": 2329,
 2024/07/17 16:27:23-07:00  "frame": 2
 2024/07/17 16:27:23-07:00 }
 2024/07/17 16:27:23-07:00 Processing frame number 2
 2024/07/17 16:27:23-07:00 Writing result
 2024/07/17 16:27:23-07:00 Waiting until a USR1 signal is sent...
 2024/07/17 16:27:23-07:00 ===
 2024/07/17 16:27:23-07:00
 2024/07/17 16:27:23-07:00 ---------------------------------------------
 2024/07/17 16:27:23-07:00 Uploading output files to Job Attachments
```

```
2024/07/17 16:27:23-07:00 ---------------------------------------------
```

The following lines from the job template are what specified this action. ``` steps:

```
steps:
- name: EnvWithDaemonProcess
  parameterSpace:
    taskParameterDefinitions:
    - name: Frame
      type: INT
      range: "{{Param.Frames}}"

  stepEnvironments:
    ...

  script:
    actions:
      onRun:
        timeout: 60
        command: bash
        args:
        - '{{Task.File.Run}}'
    embeddedFiles:
    - name: Run
      filename: run-task.sh
      type: TEXT
      data: |
        # This bash script sends a task to the background daemon process,
        # then waits for it to respond with the output result.

        set -euo pipefail

        source "$DAEMON_BASH_HELPER_SCRIPT"

        echo "Daemon PID is $DAEMON_PID"
        echo "Daemon log file is $DAEMON_LOG"

        print_daemon_log "Previous output from daemon"

        send_task_to_daemon "{\"pid\": $$, \"frame\": {{Task.Param.Frame}} }"
        wait_for_daemon_task_result

        echo Received task result:
```

```
            echo "$TASK_RESULT" | jq .

            print_daemon_log "Daemon log from running the task"

    hostRequirements:
      attributes:
      - name: attr.worker.os.family
        anyOf:
        - linux
```

Provide applications for your jobs

You can use a queue environment to load applications to process your jobs. When you create a service-managed fleet using the Deadline Cloud console, you have the option of creating a queue environment that uses the conda package manager to load applications.

If you want to use a different package manager, you can create a queue environment for that manager. For an example using Rez, see Use a different package manager.

Deadline Cloud provides a conda channel to load a selection of rendering applications into your environment. They support the submitters that Deadline Cloud provides for digital content creation applications.

You can also load software for conda-forge to use in your jobs. The following examples show job templates using the queue environment provided by Deadline Cloud to load applications before running the job.

Topics

- Getting an application from a conda channel
- Use a different package manager

Getting an application from a conda channel

You can create a custom queue environment for you Deadline Cloud workers that installs the software of your choice. This example queue environment has the same behavior as the environment used by the console for service-managed fleets. It runs conda directly to create the environment.

The environment creates a new conda virtual environment for every Deadline Cloud session that runs on a worker, and then deletes the environment when it is done.

Conda caches the downloaded packages so that they don't need to be downloaded again, but each session must link all of the packages into the environment.

The environment defines three scripts that run when Deadline Cloud starts a session on a worker. The first script runs when the onEnter action is called. It calls the other two to set up environment variables. When the script finishes running, the conda environment is available with all of the specified environment variables set.

For the latest version of the example, see conda_queue_env_console_equivalent.yaml in the deadline-cloud-samples repository on GitHub.

Process a CSV file with an application from conda-forge

The following example is a job template that process data from a CSV file using a Python package called polars. The package is loaded from the conda-forge channel.

The job sets the CondaPackages and CondaChannels parameters defined in the queue environment that tell Deadline Cloud where to get the package.

The section of the job template that sets the parameters is:

```
- name: CondaPackages
  description: A list of conda packages to install. The job expects a Queue Environment
 to handle this.
  type: STRING
  default: polars
- name: CondaChannels
  description: A list of conda channels to get packages from. The job expects a Queue
 Environment to handle this.
  type: STRING
  default: conda-forge
```

For the latest version of the complete example job template, see stage_1_self_contained_template/template.yaml. For the latest version of the queue environment that loads the conda packages, see conda_queue_env_console_equivalent.yaml in the deadline-cloud-samples repository on GitHub.

Get Blender from the deadline-cloud channel

The following example shows a job template that gets Blender from the `deadline-cloud` conda channel. This channel supports the submitters that Deadline Cloud provides for digital content creation software, though you can use the same channel to load software for your own use.

For a list of the software provided by the `deadline-cloud` channel, see [Default queue environment](#) in the *AWS Deadline Cloud Developer Guide*.

This job sets the `CondaPackages` parameter defined in the queue environment to tell Deadline Cloud to load Blender into the environment.

The section of the job template that sets the parameter is:

```
- name: CondaPackages
  type: STRING
  userInterface:
    control: LINE_EDIT
    label: Conda Packages
    groupLabel: Software Environment
  default: blender
  description: >
    Tells the queue environment to install Blender from the deadline-cloud conda
  channel.
```

For the latest version of the complete example job template, see [blender_render/template.yaml](#). For the latest version of the queue environment that loads the conda packages, see [conda_queue_env_console_equivalent.yaml](#) in the [deadline-cloud-samples](#) repository on GitHub.

Use a different package manager

The default package manager for Deadline Cloud is conda. If you need to use a different package manager, such as Rez, you can create a custom queue environment that contains scripts that use your package manager instead.

This example queue environment provides the same behavior as the environment used by the console for service-managed fleets. It replaces the conda package manager with Rez.

The environment defines three scripts that run when Deadline Cloud starts a session on a worker. The first script runs when the `onEnter` action is called. It calls the other two to set up environment

variables. When the script finishes running, the Rez environment is available with all of the specified environment variables set.

The example assumes that you are have a customer-managed fleet that uses a shared file system for the Rez packages.

For the latest version of the example, see rez_queue_env.yaml in the deadline-cloud-samples repository on GitHub.

Create a conda channel using S3

If you have custom packages for applications that are not available on the `deadline-cloud` or `conda-forge` channels you can create a conda channel that contains the packages that your environments use. You can store the packages in an Amazon S3 bucket so that you can use AWS Identity and Access Management permissions to control access to the channel.

You can also use a Deadline Cloud queue to build the packages for your conda channel to make it easier to update and maintain the application packages. The following examples show how to create a conda channel that provides Blender 4.1 for your environments.

Topics

- Create a package building queue
- Configure production queue permissions for custom conda packages
- Add a conda channel to a queue environment
- Submit the Blender 4.2 package job
- Submit a Blender 4.2 render job

Create a package building queue

In this example you create a Deadline Cloud queue to build the Blender 4.2 application. This simplifies delivery of the finished packages to the Amazon S3 bucket used as the conda channel and enables you to use your existing fleet to build the package. This reduces the number of infrastructure components to manage.

Follow the instructions in Create a queue in the *Deadline Cloud User Guide*. Make the following changes:

- In step 5, choose an existing S3 bucket. Specify a root folder name such as **DeadlineCloudPackageBuild** so that build artifacts stay separate from your normal Deadline Cloud attachments.

- In step 6, you can associate the package building queue with an existing fleet, or you can create an entirely new fleet if your current fleet is unsuitable.

- In step 9, create a new service role for your package building queue. You will modify the permissions to give the queue the permissions required for uploading packages and reindexing a conda channel.

Configure the package building queue permissions

To allow the package build queue to access the /Conda prefix in the queue's S3 bucket, you must modify the queue's role to give it read/write access. The role needs the following permissions so that package build jobs can upload new packages and reindex the channel.

- s3:GetObject
- s3:PutObject
- s3:ListBucket
- s3:GetBucketLocation
- s3:DeleteObject

1. Open the Deadline Cloud console and navigate to the queue details page for the package build queue.

2. Choose the queue service role, then choose **Edit queue**.

3. Scroll to the **Queue service role** section, then choose **View this role in the IAM console**.

4. From the list of permission policies, choose the **AmazonDeadlineCloudQueuePolicy** for your queue.

5. From the **Permissions** tab, choose **Edit**.

6. Update the queue service role to the following. Replace *amzn-s3-demo-bucket* and *111122223333* with your own bucket and account.

```
{
    "Effect": "Allow",
    "Sid": "CustomCondaChannelReadWrite",
```

```
      "Action": [
       "s3:GetObject",
       "s3:PutObject",
       "s3:DeleteObject",
       "s3:ListBucket",
       "s3:GetBucketLocation"
      ],
      "Resource": [
       "arn:aws:s3:::amzn-s3-demo-bucket",
       "arn:aws:s3:::amzn-s3-demo-bucket/Conda/*"   ],
      "Condition": {
       "StringEquals": {
        "aws:ResourceAccount": "111122223333"
       }
      }
     },
```

Configure production queue permissions for custom conda packages

Your production queue needs read-only permissions to the /Conda prefix in the queue's S3 bucket. Open the AWS Identity and Access Management (IAM) page for the role associated with the production queue and modify the policy with the following:

1. Open the Deadline Cloud console and navigate to the queue details page for the package build queue.

2. Choose the queue service role, then choose **Edit queue**.

3. Scroll to the **Queue service role** section, then choose **View this role in the IAM console**.

4. From the list of permission policies, choose the **AmazonDeadlineCloudQueuePolicy** for your queue.

5. From the **Permissions** tab, choose **Edit**.

6. Add a new section to the queue service role like the following. Replace *amzn-s3-demo-bucket* and *111122223333* with your own bucket and account.

```
{
    "Effect": "Allow",
    "Sid": "CustomCondaChannelReadOnly",
    "Action": [
     "s3:GetObject",
     "s3:ListBucket"
```

```
      ],
      "Resource": [
       "arn:aws:s3:::amzn-s3-demo-bucket",
       "arn:aws:s3:::amzn-s3-demo-bucket/Conda/*"
      ],
      "Condition": {
       "StringEquals": {
        "aws:ResourceAccount": "111122223333"
       }
      }
     },
```

Add a conda channel to a queue environment

To use the S3 conda channel, you need to add the `s3://amzn-s3-demo-bucket/Conda/`
`Default` channel location to the `CondaChannels` parameter of jobs that you submit to Deadline
Cloud. The submitters provided with Deadline Cloud provide fields to specify custom conda
channels and package.

You can avoid modifying every job by editing the conda queue environment for your production
queue. For a service-managed queue, use the following procedure:

1. Open the Deadline Cloud console and navigate to the queue details page for the production
   queue.

2. Choose the environments tab.

3. Select the **Conda** queue environment, and then choose **Edit**.

4. Choose the JSON editor, and then in the script, find the parameter definition for
   `CondaChannels`

5. Edit the line `default: "deadline-cloud"` so that it starts with the newly created S3 conda
   channel:

   ```
   default: "s3://amzn-s3-demo-bucket/Conda/Default deadline-cloud"
   ```

Service-managed fleets enable strict channel priority for conda by default, using the new S3
channel stops conda from using the `deadline-cloud` channel. Any job that successfully
completed using `blender=3.6` from the `deadline-cloud` channel will fail now that you are
using Blender 4.2.

For customer-managed fleets, you can enable the use of conda packages by using one of the conda queue environment samples in the Deadline Cloud samples GitHub repository.

Submit the Blender 4.2 package job

You can build your own Blender 4.2 conda package to render jobs, by downloading the Blender archive and then submitting a job to the package building queue. The queue sends the job to the associated fleet to build the package and reindex the conda channel.

These instructions use git from a bash-compatible shell to get an OpenJD package build job and some conda recipes from the Deadline Cloud samples GitHub repository. You also need the following:

- If you are using Windows, a version of bash, git BASH, is installed when you install git.

- You must have the Deadline Cloud CLI installed.

- You must be logged into the Deadline Cloud monitor.

1. Open the Deadline Cloud configuration GUI using the following command and set the default farm and queue to your package building queue.

   ```
   deadline config gui
   ```

2. Use the following command to clone the Deadline Cloud samples GitHub repository.

   ```
   git clone https://github.com/aws-deadline/deadline-cloud-samples.git
   ```

3. Change to the `conda_recipes` directory in the `deadline-cloud-samples` directory.

   ```
   cd deadline-cloud-samples/conda_recipes
   ```

4. Run the script called `submit-package-job`. The script provides instructions for downloading Blender the first time that you run the script.

   ```
   ./submit-package-job blender-4.2/
   ```

5. Follow the instructions for downloading Blender. When you have the archive, run the `submit-package-job` script again.

   ```
   ./submit-package-job blender-4.2/
   ```

After you submit the job, use the Deadline Cloud monitor to view the progress and status of the job as it runs.

The lower left of the monitor shows the two steps of the job, building the package and then reindexing. The lower right shows the individual steps for each task. In this example, there is one step for each task.

In the lower left of the monitor are the two steps of the job, building the package and then reindexing the conda channel. In the lower right are the individual tasks for each step. In this example there is only one task for each step.

When you right click on the task for the package building step and choose **View logs**, the monitor shows a list of session actions that show how the task is scheduled on the worker. The actions are:

- **Sync attachments** – This action copies the input job attachments or mounts a virtual file system, depending on the setting used for the job attachments file system.

- **Launch Conda** – This action is from the queue environment added by default when you created the queue. The job doesn't specify any conda packages, so it finishes quickly and doesn't create a conda virtual environment.

- **Launch CondaBuild Env** – This action creates a custom conda virtual environment that includes the software needed to build a conda package and reindex a channel. It installs from the conda-forge channel.

- **Task run** – This action builds the Blender package and uploads the results to Amazon S3.

As the actions run, they send logs in a structured format to Amazon CloudWatch. When a job is complete, select **View logs for all tasks** to see additional logs about the set up and tear down of the environment that the job runs in.

Submit a Blender 4.2 render job

After you have the Blender 4.2 package built and your production queue configured to use the S3 conda channel, you can submit jobs to render with the package. If you don't have a Blender scene, download the Blender 3.5 - Cozy Kitchen scene from the [Blender demo files](#) page.

The Deadline Cloud samples GitHub repository that you downloaded earlier contains a sample job to render a Blender scene using the following commands:

```
deadline bundle submit blender_render \
    -p CondaPackages=blender=4.2 \
    -p BlenderSceneFile=/path/to/downloaded/blender-3.5-splash.blend \
    -p Frames=1
```

You can use the Deadline Cloud monitor to track the progress of your job:

1. In the monitor, select the task for the job you submitted, then select the option to view the log.

2. On the right side of the log view, select the **Launch Conda** session action.

You can see that the action searched for Blender 4.2 in the two conda channels configured for the queue environment, and that it found the package in the S3 channel.

Create a conda package for an application

You can combine an entire application, including dependencies, into a conda package. The packages Deadline Cloud provides in the [deadline-cloud channel](#) for service-managed fleets use this binary repackaging approach. This organizes the same files as an installation to fit the conda virtual environment.

When repackaging an application for conda, there are two goals:

- Most files for the application should be separate from the primary conda virtual environment structure. Environments can then mix the application with packages from other sources like [conda-forge](#).

- When a conda virtual environment is activated, the application should be available from the PATH environment variable.

To repackage an application for conda

1. To repackage an application for conda, write conda build recipes that install the application into a subdirectory like $CONDA_PREFIX/opt/*<application-name>*. This separates it from the standard prefix directories like `bin` and `lib`.

2. Then, add symlinks or launch scripts to $CONDA_PREFIX/bin to run the application binaries.

   Alternatively, create activate.d scripts that the `conda activate` command will run to add the application binary directories to the PATH. On Windows, where symlinks are not supported everywhere environments can be created, use application launch or activate.d scripts instead.

3. Some applications depend on libraries not installed by default on Deadline Cloud service-managed fleets. For example, the X11 window system is usually unnecessary for non-interactive jobs, but some applications still require it to run without a graphical interface. You must provide those dependencies within the package you create.

4. Ensure you follow the copyright and license agreements for the applications you package. We recommend using a private Amazon S3 bucket for your conda channel to control distribution and limit package access to your farm.

Create a conda build recipe for Blender

You can use different applications to create a conda build recipe. Blender is free to use and is simple to package with conda. The Blender Foundation provides application archives for multiple operating systems. We created a sample conda build recipe that uses the Windows .zip and the Linux .tar.xz files. In this section, learn how to use the Blender 4.2 conda build recipe.

The deadline-cloud.yaml file specifies the conda platforms and other metadata for submitting package jobs to Deadline Cloud. This recipe includes local source archive information to demonstrate how that works. The linux-64 conda platform is set to build in a default job submission to match the most common configuration. The deadline-cloud.yaml looks similar to the following:

```
condaPlatforms:
  - platform: linux-64
    defaultSubmit: true
```

```
    sourceArchiveFilename: blender-4.2.1-linux-x64.tar.xz
    sourceDownloadInstructions: 'Run "curl -LO https://download.blender.org/release/
 Blender4.2/blender-4.2.1-linux-x64.tar.xz"'
  - platform: win-64
    defaultSubmit: false
    sourceArchiveFilename: blender-4.2.1-windows-x64.zip
    sourceDownloadInstructions: 'Run "curl -LO https://download.blender.org/release/
 Blender4.2/blender-4.2.1-windows-x64.zip"'
```

Review the files in the `recipe` directory. The metadata for the recipe is in recipe/meta.yaml. You can also read the conda build meta.yaml documentation to learn more, such as how the file is a template to generate YAML. The template is used to specify the version number just once, and to provide different values based on the operating system.

You can review the build options selected in `meta.yaml` to turn off various binary relocation and dynamic shared object (DSO) linking checks. These options control how the package works when it's installed into a conda virtual environment at any directory prefix. The default values simplify packaging every dependency library into a separate package, but when binary repackaging an application, you need to change them.

If the application you're packaging requires additional dependency libraries or you are packaging plugins for an application separately, you may encounter DSO errors. These errors occur when the dependency is not in the library search path for the executable or library that needs it. Applications rely on libraries being in globally defined paths, like `/lib` or `/usr/lib`, when installed on a system. However, since conda virtual environments can be placed anywhere, there is no absolute path to use. Conda uses relative RPATH features, which both Linux and macOS support, to handle this. Refer to the conda build documentation on Making packages relocatable for more information.

Blender does not require any RPATH adjustment, as the application archives were built with this in mind. For applications that do require it, you can use the same tools that conda build does: `patchelf` on Linux and `install_name_tool` on macOS.

During the package build, the build.sh or build_win.sh (called by `bld.bat`) script runs to install files into an environment prepared with the package dependencies. These scripts copy the installation files, create symlinks from `$PREFIX/bin`, and set up the activation scripts. On Windows, it does not create symlinks but instead adds the Blender directory to the PATH in the activation script.

The decision was made to use `bash` instead of a `cmd.exe` .bat file for the Windows part of the conda build recipe, as this provides more consistency across the build scripts. Refer to the Deadline Cloud developer guide's recommendation on workload portability for tips on using `bash` on Windows. If you've installed git for Windows to clone the deadline-cloud-samples git repository, you already have access to `bash`.

The conda build environment variables documentation lists the values available for use in the build script. These values include `$SRC_DIR` for the source archive data, `$PREFIX` for the installation directory, `$RECIPE_DIR` to access other files from the recipe, `$PKG_NAME` and `$PKG_VERSION` for package name and version, and `$target_platform` for the target conda platform.

Build jobs to submit to Deadline Cloud

You submit jobs to Deadline Cloud using job bundles. A job bundle is a collection of files, including an Open Job Description (OpenJD) job template and any asset files needed to render the job. The job template describes how workers process and access the assets, and provides the script that the worker runs. Job bundles enable artists, technical directors, and pipeline developers to easily submit complex jobs to Deadline Cloud from their local workstations or on-premises render farm. This is particularly useful for teams working on large-scale visual effects, animation, or other media rendering projects that require scalable, on-demand computing resources.

You can create the job bundle using the local file system to store files and a text editor to create the job template. After you create the bundle, you can use the Deadline Cloud CLI to submit the job to Deadline Cloud. Or you can use a tool, such as a Deadline Cloud submitter, to create the bundle and submit it to Deadline Cloud

You can store your assets in a file system shared between your workers, or you can use Deadline Cloud job attachments to automate moving assets to S3 buckets where your workers can access them. Job attachments also help move the output from your jobs back to your workstations.

The following sections provide detailed instructions on creating and submitting job bundles to Deadline Cloud.

Topics

- [Open Job Description (OpenJD) templates for Deadline Cloud](#)
- [Using files in your jobs](#)
- [Use job attachments to share files](#)

Open Job Description (OpenJD) templates for Deadline Cloud

A *job bundle* is one of the tools you use to define jobs for AWS Deadline Cloud. They group an Open Job Description (OpenJD) template with additional information such as files and directories that your jobs use with job attachments. You use the Deadline Cloud command-line interface to use a job bundle to submit jobs for a queue to run.

A job bundle is a directory structure that contains an OpenJD job template, other files that define the job, and job-specific files required as input for your job. You can specify the files that define your job as either YAML or JSON files.

The only required file is either `template.yaml` or `template.json`. You can also include the following files:

```
/template.yaml (or template.json)
/asset_references.yaml (or asset_references.json)
/parameter_values.yaml (or parameter_values.json)
/other job-specific files and directories
```

Use a job bundle for custom job submissions with the Deadline Cloud CLI and a job attachment, or you can use an graphical submission interface. For example, the following is the Blender sample from GitHub. You can run the sample using this command in the Blender sample directory:

```
deadline bundle gui-submit blender_render
```
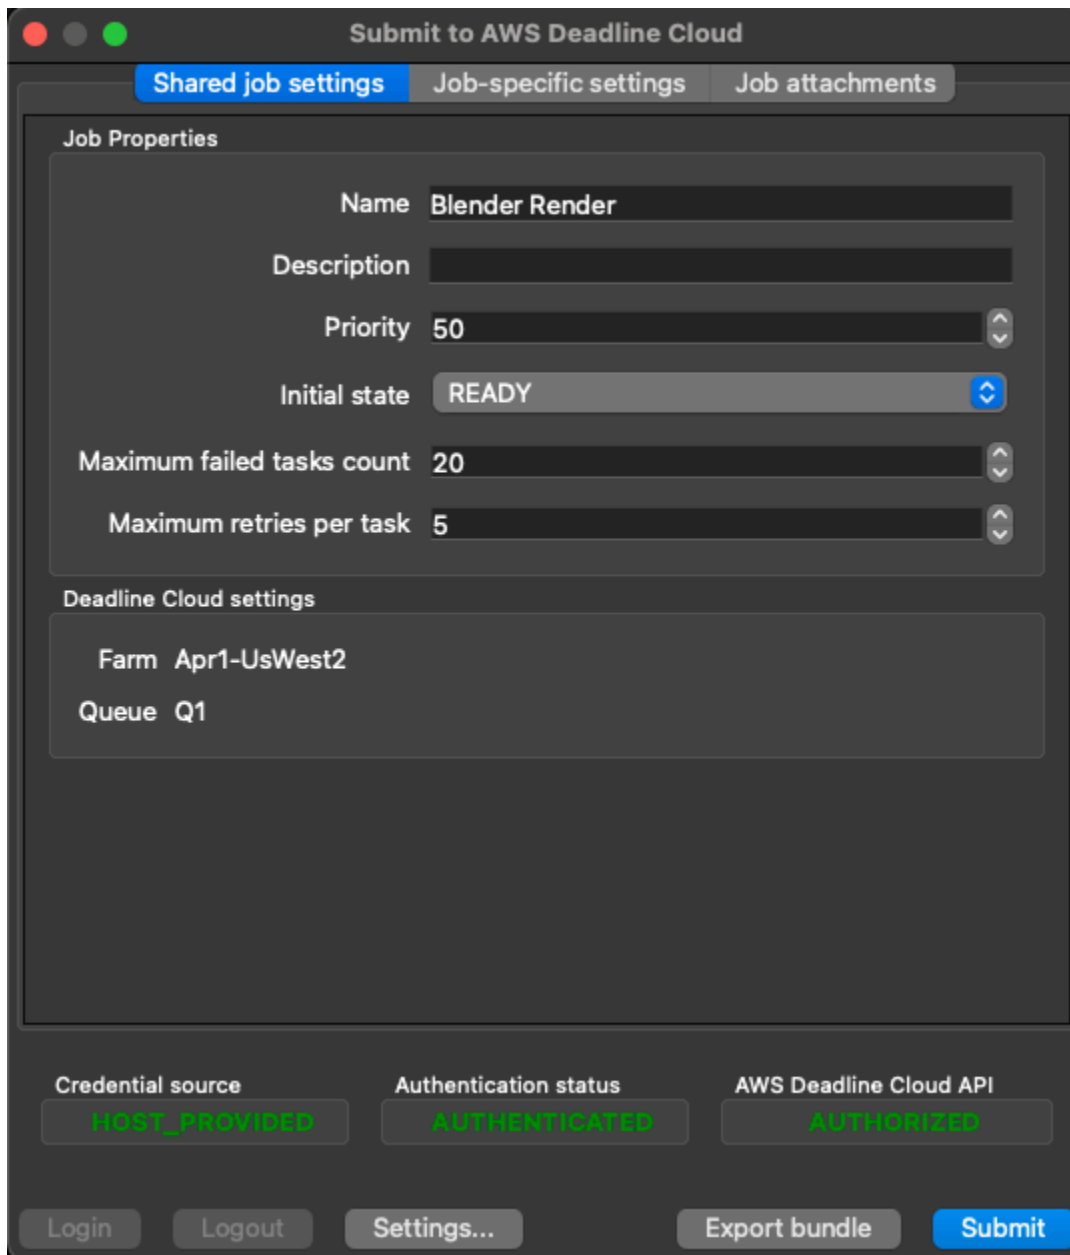
The job-specific settings panel are generated from the `userInterface` properties of the job parameters defined in the job template.

To submit a job using the command line, you can use a command similar to the following

```
deadline bundle submit \
    --yes \
    --name Demo \
     -p BlenderSceneFile=location of scene file \
     -p OutputDir=file pathe for job output \
```

```
        blender_render/
```

Or you can use the `deadline.client.api.create_job_from_job_bundle` function in the `deadline` Python package.

All of the job submitter plugins provided with Deadline Cloud, such as the Autodesk Maya plugin, generate a job bundle for your submission and then use the Deadline Cloud Python package to submit your job to Deadline Cloud. You can see the job bundles submitted in the job history directory of you workstation or by using a submitter. You can find your job history directory with the following command:

```
deadline config get settings.job_history_dir
```

When you job is running on a Deadline Cloud worker, it has access to environment variables that provide it with information about the job. The environment variables are:

Variable name	Available
DEADLINE_FARM_ID	All actions
DEADLINE_FLEET_ID	All actions
DEADLINE_WORKER_ID	All actions
DEADLINE_QUEUE_ID	All actions
DEADLINE_JOB_ID	All actions
DEADLINE_SESSION_ID	All actions
DEADLINE_SESSIONACTION_ID	All actions
DEADLINE_TASK_ID	Task actions

Topics

- Job template elements for job bundles
- Parameter values elements for job bundles
- Asset references elements for job bundles

Job template elements for job bundles

The job template defines the runtime environment and the processes that run as part of a Deadline Cloud job. You can create parameters in a template so that it can be used to create jobs that differ only in input values, much like a function in a programming language.

When you submit a job to Deadline Cloud, it runs in any queue environments applied to the queue. Queue environments are built using the Open Job Description (OpenJD) external environments specification. For details, see the Environment template in the OpenJD GitHub repository.

For an introduction creating a job with an OpenJD job template, see Introduction to creating a job in the OpenJD GitHub repository. Additional information can be found in How jobs are run. There are job template samples in the in the OpenJD GitHub repository's `samples` directory.

You can define the job template in either YAML format (`template.yaml`) or JSON format (`template.json`). The examples in this section are shown in YAML format.

For example, the job template for the `blender_render` sample defines an input parameter `BlenderSceneFile` as a file path:

```
- name: BlenderSceneFile
  type: PATH
  objectType: FILE
  dataFlow: IN
  userInterface:
    control: CHOOSE_INPUT_FILE
    label: Blender Scene File
    groupLabel: Render Parameters
    fileFilters:
    - label: Blender Scene Files
      patterns: ["*.blend"]
    - label: All Files
      patterns: ["*"]
  description: >
    Choose the Blender scene file to render. Use the 'Job Attachments' tab
    to add textures and other files that the job needs.
```

The `userInterface` property defines the behavior of automatically generated user interfaces for both the command line using the `deadline bundle gui-submit` command and within the job submission plugins for applications like Autodesk Maya.

In this example, the UI widget for inputting a value for the `BlenderSceneFile` parameter is a file-selection dialog that shows only `.blend` files.

For more examples of using the `userInteface` element, see the [gui_control_showcase](#) sample in the [deadline-cloud-samples](#) repository on GitHub.

The `objectType` and `dataFlow` properties control the behavior of job attachments when you submit a job from a job bundle. In this case, `objectType: FILE` and `dataFlow:IN` mean that the value of `BlenderSceneFile` is an input file for job attachments.

In contrast, the definition of the `OutputDir` parameter has `objectType: DIRECTORY` and `dataFlow: OUT`:

```
- name: OutputDir
  type: PATH
  objectType: DIRECTORY
  dataFlow: OUT
  userInterface:
    control: CHOOSE_DIRECTORY
    label: Output Directory
    groupLabel: Render Parameters
  default: "./output"
  description: Choose the render output directory.
```

The value of the `OutputDir` parameter is used by job attachments as the directory where the job writes output files.

For more information about the `objectType` and `dataFlow` properties, see [JobPathParameterDefinition](#) in the [Open Job Description specification](#)

The rest of the `blender_render` job template sample defines the job's workflow as a singe step with each frame in the animation rendered as a separate task:

```
steps:
- name: RenderBlender
  parameterSpace:
    taskParameterDefinitions:
```

```
    - name: Frame
      type: INT
      range: "{{Param.Frames}}"
  script:
    actions:
      onRun:
        command: bash
        # Note: {{Task.File.Run}} is a variable that expands to the filename on the
 worker host's
        # disk where the contents of the 'Run' embedded file, below, is written.
        args: ['{{Task.File.Run}}']
    embeddedFiles:
      - name: Run
        type: TEXT
        data: |
          # Configure the task to fail if any individual command fails.
          set -xeuo pipefail

          mkdir -p '{{Param.OutputDir}}'

          blender --background '{{Param.BlenderSceneFile}}' \
                  --render-output '{{Param.OutputDir}}/{{Param.OutputPattern}}' \
                  --render-format {{Param.Format}} \
                  --use-extension 1 \
                  --render-frame {{Task.Param.Frame}}
```

For example, if the value of the `Frames` parameter is 1-10, it defines 10 tasks. Each has task has a different value for the `Frame` parameter. To run a task:

1. All of the variable references in the `data` property of the embedded file are expanded, for example `--render-frame 1`.

2. The contents of the `data` property is written to a file in the session working directory on disk.

3. The task's onRun command resolves to bash *location of embedded file* and then runs.

For more information about embedded files, sessions, and path-mapped locations, see How jobs are run in the Open Job Description specification.

There are more examples of job templates in the deadline-cloud-samples/job_bundles repository, as well as the template samples provided with the Open Job Descriptions specification.

Parameter values elements for job bundles

You can use the parameters file to set the values of some of the job parameters in the job template or [CreateJob](#) operation request arguments in the job bundle so that you don't need to set values when submitting a job. The UI for job submission enables you to modify these values.

You can define the job template in either YAML format (`parameter_values.yaml`) or JSON format (`parameter_values.json`). The examples in this section are shown in YAML format.

In YAML, the format of the file is:

```
parameterValues:
- name: <string>
  value: <integer>, <float>, or <string>
- name: <string>
  value: <integer>, <float>, or <string>ab
... repeating as necessary
```

Each element of the `parameterValues` list must be one of the following:

- A job parameter defined in the job template.

- A job parameter defined in a queue environment for the queue that you submit the job to..

- A special parameter passed to the `CreateJob` operation when creating a job.

  - `deadline:priority` – The value must be an integer. It is passed to the `CreateJob` operation as the [priority](#) parameter.

  - `deadline:targetTaskRunStatus` – The value must be a string. It is passed to the `CreateJob` operation as the [targetTaskRunStatus](#) parameter.

  - `deadline:maxFailedTasksCount` – The value must be an integer. It is passed to the `CreateJob` operation as the [priority](#) parameter.

  - `deadline:maxRetriesPerTask` – The value must be an integer. It is passed to the `CreateJob` operation as the [priority](#) parameter.

A job template is always a template rather than a specific job to run. A parameter values file enables a job bundle to either act as a template if some parameters don't have values defined in this file, or as a specific job submission if all parameters have values.

For example, the [blender_render sample](#) doesn't have a parameters file and its job template defines parameters with no default values. This template must be used as a template to create

jobs. After you create a job using this job bundle, Deadline Cloud writes a new job bundle to the job history directory.

For example, when you submit a job with the following command:

```
deadline bundle gui-submit blender_render/
```

The new job bundle contains a `parameter_values.yaml` file that contains the specified parameters:

```
% cat ~/.deadline/job_history/\(default\)/2024-06/2024-06-20-01-JobBundle-Demo/
parameter_values.yaml
parameterValues:
- name: deadline:targetTaskRunStatus
  value: READY
- name: deadline:maxFailedTasksCount
  value: 10
- name: deadline:maxRetriesPerTask
  value: 5
- name: deadline:priority
  value: 75
- name: BlenderSceneFile
  value: /private/tmp/bundle_demo/bmw27_cpu.blend
- name: Frames
  value: 1-10
- name: OutputDir
  value: /private/tmp/bundle_demo/output
- name: OutputPattern
  value: output_####
- name: Format
  value: PNG
- name: CondaPackages
  value: blender
- name: RezPackages
  value: blender
```

You can create the same job with the following command:

```
deadline bundle submit ~/.deadline/job_history/\(default\)/2024-06/2024-06-20-01-
JobBundle-Demo/
```

> **ⓘ Note**
>
> The job bundle that you submit is saved to your job history directory. You can find the location of that directory with the following command:
>
> ```
> deadline config get settings.job_history_dir
> ```

Asset references elements for job bundles

You can use Deadline Cloud job attachments to transfer files back and forth between your workstation and Deadline Cloud. The asset reference file lists input files and directories, as well as output directories for your attachments. If you don't list all of the files and directories in this file, you can select them when you submit a job with the `deadline bundle gui-submit` command.

This file has no effect if you are not using job attachments.

You can define the job template in either YAML format (`asset_references.yaml`) or JSON format (`asset_references.json`). The examples in this section are shown in YAML format.

In YAML, the format of the file is:

```
assetReferences:
    inputs:
        # Filenames on the submitting workstation whose file contents are needed as
        # inputs to run the job.
        filenames:
        - list of file paths
        # Directories on the submitting workstation whose contents are needed as inputs
        # to run the job.
        directories:
        - list of directory paths

    outputs:
        # Directories on the submitting workstation where the job writes output files
        # if running locally.
        directories:
        - list of directory paths

    # Paths referenced by the job, but not necessarily input or output.
```

```
    # Use this if your job uses the name of a path in some way, but does not explicitly
need
    # the contents of that path.
    referencedPaths:
    - list of directory paths
```

When selecting the input or output file to upload to Amazon S3, Deadline Cloud compares the file path against the paths listed in your storage profiles. Each SHARED-type file system location in a storage profile abstracts a network file share that is mounted on your workstations and worker hosts. Deadline Cloud uploads only files that are not on one of these file shares.

For more information about creating and using storage profiles, see Shared storage in Deadline Cloud in the *AWS Deadline Cloud User Guide*.

Example - The asset reference file created by the Deadline Cloud GUI

Use the following command to submit a job using the blender_render sample.

```
deadline bundle gui-submit blender_render/
```

Add some additional files to the job on the **Job attachments** tab:

After you submit the job, you can look at the `asset_references.yaml` file in the job bundle in the job history directory to see the assets in the YAML file:

```
% cat ~/.deadline/job_history/\(default\)/2024-06/2024-06-20-01-JobBundle-Demo/
asset_references.yaml
```

```
assetReferences:
  inputs:
    filenames:
    - /private/tmp/bundle_demo/a_texture.png
    directories:
    - /private/tmp/bundle_demo/assets
  outputs:
    directories: []
  referencedPaths: []
```

Using files in your jobs

Many of the jobs that you submit to AWS Deadline Cloud have input and output files. Your input files and output directories may be located on a combination of shared filesystems and local drives. Jobs need to locate the content in those locations. Deadline Cloud provides two features, job attachments and storage profiles that work together to help your jobs locate the files that they need.

Job attachments helps you move files to your worker hosts from filesystem locations on your workstation that are not available on your worker hosts, and vice versa. It moves files between hosts using Amazon S3. You can enable job attachments on each of your queues to make it available to jobs in those queues. Job attachments are used primarily with service-managed fleets, but you can also use them with customer-managed fleets.

Use storage profiles to model the layout of shared filesystem locations on your workstation and worker hosts. This helps your jobs locate shared files and directories when their locations differ between your workstation and worker hosts, such as cross-platform setups with Windows-based workstations and Linux-based worker hosts. Storage profile's model of your filesystem configuration is also used by job attachments to identify the files it needs to shuttle between hosts through Amazon S3.

If you are not using job attachments, and you don't need to remap file and directory locations between workstations and worker hosts then you don't need to model your fileshares with storage profiles.

Topics

- Sample project infrastructure
- Storage profiles and path mapping

Sample project infrastructure

To demonstrate using job attachments and storage profiles, set up a test environment with two separate projects. You can use the Deadline Cloud console to create the test resources.

1. If you haven't already, create a test farm. To create a farm, follow the procedure in Create a farm.

2. Create two queues for jobs in each of the two projects. To create queues, follow the procedure in Create a queue.

   a. Create the first queue called **Q1**. Use the following configuration, use the defaults for all other items.

      - For job attachments, choose **Create a new Amazon S3 bucket**.

      - Select **Enable association with customer-managed fleets**.

      - For the run as user, enter **jobuser** for both the POSIX user and group.

      - For the queue service role, create a new role named **AssetDemoFarm-Q1-Role**

      - Clear the default Conda queue environment checkbox.

   b. Create the second queue called **Q2**. Use the following configuration, use the defaults for all other items.

      - For job attachments, choose **Create a new Amazon S3 bucket**.

      - Select **Enable association with customer-managed fleets**.

      - For the run as user, enter **jobuser** for both the POSIX user and group.

      - For the queue service role, create a new role named **AssetDemoFarm-Q2-Role**

      - Clear the default Conda queue environment checkbox.

3. Create a single customer-managed fleet that runs the jobs from both queues. To create the fleet, follow the procedure in Create a customer-managed fleet. Use the following configuration:

   - For **Name**, use **DemoFleet**.

   - For **Fleet type** choose **Customer managed**

   - For **Fleet service role**, create a new role named **AssetDemoFarm-Fleet-Role**.

   - Don't associate the fleet with any queues.

The test environment assumes that there are three file systems shared between hosts using network file shares. In this example, the locations have the following names:

- `FSCommon` - contains input job assets that are common to both projects.

- `FS1` - contains input and output job assets for project 1.

- `FS2` - contains input and output job assets for project 2.

The test environment also assumes that there are three workstations, as follows:

- `WSAll` - A Linux-based workstation used by developers for all projects. The shared file system locations are:

  - FSCommon: `/shared/common`

  - FS1: `/shared/projects/project1`

  - FS2: `/shared/projects/project2`

- `WS1` - A Windows-based workstation used for project 1. The shared file system locations are:

  - FSCommon: `S:\`

  - FS1: `Z:\`

  - FS2: Not available

- `WS1` - A macOS-based workstation used for project 2.The shared file system locations are:

  - FSCommon: `/Volumes/common`

  - FS1: Not available

  - FS2: `/Volumes/projects/project2`

Finally, define the shared file system locations for the workers in your fleet. The examples that follow refer to this configuration as `WorkerConfig`. The shared locations are:

- FSCommon: `/mnt/common`

- FS1: `/mnt/projects/project1`

- FS2: `/mnt/projects/project2`

You don't need to set up any shared file systems, workstations, or workers that match this configuration. The shared locations don't need to exist for the demonstration.

Storage profiles and path mapping

Use storage profiles to model the file systems on your workstation and worker hosts. Each storage profile describes the operating system and file system layout of one of your system configurations. This topic describes how to use storage profiles to model the file system configurations of your hosts so Deadline Cloud can generate path mapping rules for your jobs, and how those path mapping rules are generated from your storage profiles.

When you submit a job to Deadline Cloud you can provide an optional storage profile ID for the job. This storage profile describes the submitting workstation's file system. It describes the original file system configuration that the file paths in the job template use.

You can also associate a storage profile with a [customer-managed fleet](customer-managed-fleet). The storage profile describes the file system configuration of all worker hosts in the fleet. If you have workers with different file system configuration, those workers must be assigned to a different fleet in your farm. Storage profiles are not supported in [service-managed fleets](service-managed-fleets).

Path mapping rules describe how paths should be remapped from how they are specified in the job to the path's actual location on a worker host. Deadline Cloud compares the file system configuration described in a job's storage profile with the storage profile of the fleet that is running the job to derive these path mapping rules.

Topics

- [Model shared file system locations with storage profiles](model-shared-file-system-locations)
- [Configure storage profiles for fleets](configure-storage-profiles-for-fleets)
- [Configure storage profiles for queues](configure-storage-profiles-for-queues)
- [Derive path mapping rules from storage profiles](derive-path-mapping-rules)

Model shared file system locations with storage profiles

A storage profile models the file system configuration of one of your host configurations. There are four different host configurations in the [sample project infrastructure](sample-project-infrastructure). In this example you create a separate storage profile for each. You can create a storage profile using any of the following:

- [CreateStorageProfile API](create-storage-profile-api)
- [AWS::Deadline::StorageProfile](aws-deadline-storageprofile) AWS CloudFormation resource
- [AWS console](aws-console)

A storage profile is made up of a list of file system locations that each tell Deadline Cloud the location and type of a file system location that is relevant for jobs submitted from or run on a host. A storage profile should only model the locations that are relevant for jobs. For example, the shared FSCommon location is located on workstation WS1 at `S:\`, so the corresponding file system location is:

```
{
    "name": "FSCommon",
    "path": "S:\\",
    "type": "SHARED"
}
```

Use the following commands to create the storage profile for workstation configurations WS1, WS2, and WS3 and the worker configuration `WorkerConfig` using the [AWS CLI](#) in [AWS CloudShell](#):

```
# Change the value of FARM_ID to your farm's identifier
FARM_ID=farm-00112233445566778899aabbccddeeff

aws deadline create-storage-profile --farm-id $FARM_ID \
  --display-name WSAll \
  --os-family LINUX \
  --file-system-locations \
  '[
      {"name": "FSCommon", "type":"SHARED", "path":"/shared/common"},
      {"name": "FS1", "type":"SHARED", "path":"/shared/projects/project1"},
      {"name": "FS2", "type":"SHARED", "path":"/shared/projects/project2"}
  ]'

aws deadline create-storage-profile --farm-id $FARM_ID \
  --display-name WS1 \
  --os-family WINDOWS \
  --file-system-locations \
  '[
      {"name": "FSCommon", "type":"SHARED", "path":"S:\\"},
      {"name": "FS1", "type":"SHARED", "path":"Z:\\"}
   ]'

aws deadline create-storage-profile --farm-id $FARM_ID \
  --display-name WS2 \
  --os-family MACOS \
  --file-system-locations \
  '[
```

```
        {"name": "FSCommon", "type":"SHARED", "path":"/Volumes/common"},
        {"name": "FS2", "type":"SHARED", "path":"/Volumes/projects/project2"}
    ]'

aws deadline create-storage-profile --farm-id $FARM_ID \
    --display-name WorkerCfg \
    --os-family LINUX \
    --file-system-locations \
    '[
        {"name": "FSCommon", "type":"SHARED", "path":"/mnt/common"},
        {"name": "FS1", "type":"SHARED", "path":"/mnt/projects/project1"},
        {"name": "FS2", "type":"SHARED", "path":"/mnt/projects/project2"}
    ]'
```

> **ⓘ Note**
>
> You must refer to the file system locations in your storage profiles using the same values for the name property across all storage profiles in your farm. Deadline Cloud compares the names to determine that file system locations from different storage profiles are referring to the same location when generating path mapping rules.

Configure storage profiles for fleets

The configuration of a customer-managed fleet can include a storage profile that models the file system locations on all workers in the fleet. The host file system configuration of all workers in a fleet must match their fleet's storage profile. Workers with different file system configurations must be in separate fleets.

To set your fleet's configuration to use the `WorkerConfig` storage profile use the [AWS CLI](#) in [AWS CloudShell](#):

```
# Change the value of FARM_ID to your farm's identifier
FARM_ID=farm-00112233445566778899aabbccddeeff
# Change the value of FLEET_ID to your fleet's identifier
FLEET_ID=fleet-00112233445566778899aabbccddeeff
# Change the value of WORKER_CFG_ID to your storage profile named WorkerConfig
WORKER_CFG_ID=sp-00112233445566778899aabbccddeeff

FLEET_WORKER_MODE=$( \
    aws deadline get-fleet --farm-id $FARM_ID --fleet-id $FLEET_ID \
```

```
    --query '.configuration.customerManaged.mode' \
)
FLEET_WORKER_CAPABILITIES=$( \
  aws deadline get-fleet --farm-id $FARM_ID --fleet-id $FLEET_ID \
    --query '.configuration.customerManaged.workerCapabilities' \
)

aws deadline update-fleet --farm-id $FARM_ID --fleet-id $FLEET_ID \
  --configuration \
  "{
    \"customerManaged\": {
      \"storageProfileId\": \"$WORKER_CFG_ID\",
      \"mode\": $FLEET_WORKER_MODE,
      \"workerCapabilities\": $FLEET_WORKER_CAPABILITIES
    }
  }"
```

Configure storage profiles for queues

A queue's configuration includes a list of case-sensitive names of the shared file system locations that jobs submitted to the queue require access to. for example, jobs submitted to queue Q1 require file system locations FSCommon and FS1. Jobs submitted to queue Q2 require file system locations FSCommon and FS2.

To set the queue's configurations to require these file system locations, use the following script:

```
# Change the value of FARM_ID to your farm's identifier
FARM_ID=farm-00112233445566778899aabbccddeeff
# Change the value of QUEUE1_ID to queue Q1's identifier
QUEUE1_ID=queue-00112233445566778899aabbccddeeff
# Change the value of QUEUE2_ID to queue Q2's identifier
QUEUE2_ID=queue-00112233445566778899aabbccddeeff

aws deadline update-queue --farm-id $FARM_ID --queue-id $QUEUE1_ID \
  --required-file-system-location-names-to-add FSComm FS1

aws deadline update-queue --farm-id $FARM_ID --queue-id $QUEUE2_ID \
  --required-file-system-location-names-to-add FSComm FS2
```

> **ⓘ Note**
>
> If a queue has any required files system locations, that queue can't be associated with a service-managed fleet because the fleet can't mount your shared file systems.

A queue's configuration also includes a list of allowed storage profiles that applies to jobs submitted to and fleets associated with that queue. Only storage profiles that define file system locations for all of the required file system locations for the queue are allowed in the queue's list of allowed storage profiles.

A job fails if you submit it with a storage profile that isn't in the list of allowed storage profiles for the queue. You can always submit a job with no storage profile to a queue. The workstation configurations labeled `WSAll` and `WS1` both have the required file system locations (`FSCommon` and `FS1`) for queue Q1. They need to be allowed to submit jobs to the queue. Similarly, workstation configurations `WSAll` and `WS2` meet the requirements for queue Q2. They need to be allowed to submit jobs to that queue. Update both queue configurations to allow jobs to be submitted with these storage profiles using the following script:

```
# Change the value of WSALL_ID to the identifier of the WSAll storage profile
WSALL_ID=sp-00112233445566778899aabbccddeeff
# Change the value of WS1 to the identifier of the WS1 storage profile
WS1_ID=sp-00112233445566778899aabbccddeeff
# Change the value of WS2 to the identifier of the WS2 storage profile
WS2_ID=sp-00112233445566778899aabbccddeeff

aws deadline update-queue --farm-id $FARM_ID --queue-id $QUEUE1_ID \
  --allowed-storage-profile-ids-to-add $WSALL_ID $WS1_ID

aws deadline update-queue --farm-id $FARM_ID --queue-id $QUEUE2_ID \
  --allowed-storage-profile-ids-to-add $WSALL_ID $WS2_ID
```

If you add the `WS2` storage profile to the list of allowed storage profiles for queue Q1 it fails:

```
$ aws deadline update-queue --farm-id $FARM_ID --queue-id $QUEUE1_ID \
  --allowed-storage-profile-ids-to-add $WS2_ID

An error occurred (ValidationException) when calling the UpdateQueue operation: Storage
 profile id: sp-00112233445566778899aabbccddeeff does not have required file system
 location: FS1
```

This is because the WS2 storage profile doesn't contain a definition for the file system location named FS1 that queue Q1 requires.

Associating a fleet that is configured with a storage profile that is not in the queue's list of allowed storage profiles also fails. For example:

```
$ aws deadline create-queue-fleet-association --farm-id $FARM_ID \
    --fleet-id $FLEET_ID \
    --queue-id $QUEUE1_ID

An error occurred (ValidationException) when calling the CreateQueueFleetAssociation
 operation: Mismatch between storage profile ids.
```

To fix the error, add the storage profile named `WorkerConfig` to the list of allowed storage profiles for both queue Q1 and queue Q2. Then, associate the fleet with these queues so that workers in the fleet can run jobs from both queues.

```
# Change the value of FLEET_ID to your fleet's identifier
FLEET_ID=fleet-00112233445566778899aabbccddeeff
# Change the value of WORKER_CFG_ID to your storage profile named WorkerCfg
WORKER_CFG_ID=sp-00112233445566778899aabbccddeeff

aws deadline update-queue --farm-id $FARM_ID --queue-id $QUEUE1_ID \
    --allowed-storage-profile-ids-to-add $WORKER_CFG_ID

aws deadline update-queue --farm-id $FARM_ID --queue-id $QUEUE2_ID \
    --allowed-storage-profile-ids-to-add $WORKER_CFG_ID

aws deadline create-queue-fleet-association --farm-id $FARM_ID \
    --fleet-id $FLEET_ID \
    --queue-id $QUEUE1_ID

aws deadline create-queue-fleet-association --farm-id $FARM_ID \
    --fleet-id $FLEET_ID \
    --queue-id $QUEUE2_ID
```

Derive path mapping rules from storage profiles

Path mapping rules describe how paths should be remapped from the job to the path's actual location on a worker host. When a task is running on a worker, the storage profile from the job is compared to the storage profile of the worker's fleet to derive the path mapping rules for the task.

Deadline Cloud creates a mapping rule for each of the required file system locations in the queue's configuration. For example, a job submitted with the WSAll storage profile to queue Q1 has the path mapping rules:

- FSComm: /shared/common -> /mnt/common

- FS1: /shared/projects/project1 -> /mnt/projects/project1

Deadline Cloud creates rules for the FSComm and FS1 file system locations, but not the FS2 file system location even though both the WSAll and WorkerConfig storage profiles define FS2. This is because queue Q1's list of required file system locations is ["FSComm", "FS1"].

You can confirm the path mapping rules available to jobs submitted with a particular storage profile by submitting a job that prints out [Open Job Description's path mapping rules file](#), and then reading the session log after the job has completed:

```
# Change the value of FARM_ID to your farm's identifier
FARM_ID=farm-00112233445566778899aabbccddeeff
# Change the value of QUEUE1_ID to queue Q1's identifier
QUEUE1_ID=queue-00112233445566778899aabbccddeeff
# Change the value of WSALL_ID to the identifier of the WSALL storage profile
WSALL_ID=sp-00112233445566778899aabbccddeeff

aws deadline create-job --farm-id $FARM_ID --queue-id $QUEUE1_ID \
  --priority 50 \\
  --storage-profile-id $WSALL_ID \
  --template-type JSON --template \
  '{
    "specificationVersion": "jobtemplate-2023-09",
    "name": "DemoPathMapping",
    "steps": [
      {
        "name": "ShowPathMappingRules",
        "script": {
          "actions": {
            "onRun": {
              "command": "/bin/cat",
              "args": [ "{{Session.PathMappingRulesFile}}" ]
            }
          }
        }
      }
```

```
    ]
  }'
```

If you use the [Deadline Cloud CLI](#) to submit jobs, its configuration
`settings.storage_profile_id` setting sets the storage profile that jobs submitted with the
CLI will have. To submit jobs with the `WSAll` storage profile, set:

```
deadline config set settings.storage_profile_id $WSALL_ID
```

To run a customer-managed worker as though it is running in the sample infrastructure, follow
the procedure in [Run the worker agent](#) in the *Deadline Cloud User Guide* to run a worker with
AWS CloudShell. If you followed those instructions before, delete the ~/demoenv-logs and ~/
demoenv-persist directories first. Also, set the values of the DEV_FARM_ID and DEV_CMF_ID
environment variables that the directions reference as follows before doing so:

```
DEV_FARM_ID=$FARM_ID
DEV_CMF_ID=$FLEET_ID
```

After the job runs, you can see the path mapping rules in the job's log file:

```
cat demoenv-logs/${QUEUE1_ID}/*.log
...
JJSON log results (see below)
...
```

The log contains mapping for both the FS1 and FSComm file systems. Reformatted for readability,
the log entry looks like this:

```
{
    "version": "pathmapping-1.0",
    "path_mapping_rules": [
        {
            "source_path_format": "POSIX",
            "source_path": "/shared/projects/project1",
            "destination_path": "/mnt/projects/project1"
        },
        {
            "source_path_format": "POSIX",
            "source_path": "/shared/common",
            "destination_path": "/mnt/common"
```

```
        }
    ]
```

You can submit jobs with different storage profiles to see how the path mapping rules change.

Use job attachments to share files

Use *job attachments* to make files not in shared directories available for your jobs, and to capture the output files if they are not written to shared directories. Job attachments uses Amazon S3 to shuttle files between hosts. Files are stored in S3 buckets, and you don't need to upload a file if its content hasn't changed.

You must use job attachments when running jobs on service-managed fleets because hosts don't share file system locations. Job attachments are also useful with customer-managed fleets when a job's input or output files stored on a shared network file system, such as when your job bundle contains shell or Python scripts.

When you submit a job bundle with either the Deadline Cloud CLI or a Deadline Cloud submitter, job attachments use the job's storage profile and the queue's required file system locations to identify the input files that are not on a worker host and should be uploaded to Amazon S3 as part of job submission. These storage profiles also help Deadline Cloud identify the output files in worker host locations that must be uploaded to Amazon S3 so that they are available to your workstation.

The job attachments examples use the farm, fleet, queues, and storage profiles configurations from Sample project infrastructure and Storage profiles and path mapping. You should go through those sections before this one.

In the following examples, you use a sample job bundle as a starting point, then modify it to explore job attachment's functionality. Job bundles are the best way for your jobs to use job attachments. They combine an Open Job Description job template in a directory with additional files that list the files and directories required by jobs using the job bundle. For more information about job bundles, see Open Job Description (OpenJD) templates for Deadline Cloud.

Submitting files with a job

With Deadline Cloud, you can enable job workflows to access input files that are unavailable in shared file system locations on worker hosts. Job attachments allow rendering jobs to access files residing only on a local workstation drive or a service-managed fleet environment. When

submitting a job bundle, you can include lists of input files and directories required by the job. Deadline Cloud identifies these non-shared files, uploads them from the local machine to Amazon S3, and downloads them to the worker host. It streamlines the process of transferring input assets to render nodes, ensuring all required files are accessible for distributed job execution.

You can specify the files for jobs directly in the job bundle, use parameters in the job template that you provide using environment variables or a script, and use the job's `assets_references` file. You can use one of these methods or a combination of all three. You can specify a storage profile for the bundle for the job so that it only uploads files that have changed on the local workstation.

This section uses an example job bundle from GitHub to demonstrate how Deadline Cloud identifies the files in your job to upload, how those files are organized in Amazon S3, and how they are made available to the worker hosts processing your jobs.

Topics

- [How Deadline Cloud uploads files to Amazon S3](#)
- [How Deadline Cloud chooses the files to upload](#)
- [How jobs find job attachment input files](#)

How Deadline Cloud uploads files to Amazon S3

This example shows how Deadline Cloud uploads files from your workstation or worker host to Amazon S3 so that they can be shared. It uses a sample job bundle from GitHub and the Deadline Cloud CLI to submit jobs.

Start by cloning the [Deadline Cloud samples GitHub repository](#) into your [AWS CloudShell](#) environment, then copy the `job_attachments_devguide` job bundle into your home directory:

```
git clone https://github.com/aws-deadline/deadline-cloud-samples.git
cp -r deadline-cloud-samples/job_bundles/job_attachments_devguide ~/
```

Install the [Deadline Cloud CLI](#) to submit job bundles:

```
pip install deadline --upgrade
```

The `job_attachments_devguide` job bundle has a single step with a task that runs a bash shell script whose file system location is passed as a job parameter. The job parameter's definition is:

```
...
```

```
  - name: ScriptFile
    type: PATH
    default: script.sh
    dataFlow: IN
    objectType: FILE
...
```

The `dataFlow` property's `IN` value tells job attachments that the value of the `ScriptFile` parameter is an input to the job. The value of the `default` property is a relative location to the job bundle's directory, but it can also be an absolute path. This parameter definition declares the `script.sh` file in the job bundle's directory as an input file required for the job to run.

Next, make sure that the Deadline Cloud CLI does not have a storage profile configured then submit the job to queue Q1:

```
# Change the value of FARM_ID to your farm's identifier
FARM_ID=farm-00112233445566778899aabbccddeeff
# Change the value of QUEUE1_ID to queue Q1's identifier
QUEUE1_ID=queue-00112233445566778899aabbccddeeff

deadline config set settings.storage_profile_id ''

deadline bundle submit --farm-id $FARM_ID --queue-id $QUEUE1_ID
  job_attachments_devguide/
```

The output from the Deadline Cloud CLI after this command is run looks like:

```
Submitting to Queue: Q1
...
Hashing Attachments   [################################]   100%
Hashing Summary:
    Processed 1 file totaling 39.0 B.
    Skipped re-processing 0 files totaling 0.0 B.
    Total processing time of 0.0327 seconds at 1.19 KB/s.

Uploading Attachments   [###############################]   100%
Upload Summary:
    Processed 1 file totaling 39.0 B.
    Skipped re-processing 0 files totaling 0.0 B.
    Total processing time of 0.25639 seconds at 152.0 B/s.

Waiting for Job to be created...
```

```
Submitted job bundle:
    job_attachments_devguide/
Job creation completed successfully
job-74148c13342e4514b63c7a7518657005
```

When you submit the job, Deadline Cloud first hashes the `script.sh` file and then it uploads it to Amazon S3.

Deadline Cloud treats the S3 bucket as content-addressable storage. Files are uploaded to S3 objects. The object name is derived from a hash of the file's contents. If two files have identical contents they have the same hash value regardless of where the files are located or what they are named. This enables Deadline Cloud to avoid uploading a file if it is already available.

You can use the [AWS CLI](#) to see the objects that were uploaded to Amazon S3:

```
# The name of queue `Q1`'s job attachments S3 bucket
Q1_S3_BUCKET=$(
  aws deadline get-queue --farm-id $FARM_ID --queue-id $QUEUE1_ID \
    --query 'jobAttachmentSettings.s3BucketName' | tr -d '"'
)

aws s3 ls s3://$Q1_S3_BUCKET --recursive
```

Two objects were uploaded to S3:

- `DeadlineCloud/Data/87cb19095dd5d78fcaf56384ef0e6241.xxh128` – The contents of `script.sh`. The value `87cb19095dd5d78fcaf56384ef0e6241` in the object key is the hash of the file's contents, and the extension `xxh128` indicates that the hash value was calculated as a 128 bit [xxhash](#).

- `DeadlineCloud/Manifests/<farm-id>/<queue-id>/Inputs/<guid>/a1d221c7fd97b08175b3872a37428e8c_input` – The manifest object for the job submission. The values `<farm-id>`, `<queue-id>`, and `<guid>` are your farm identifier, queue identifier, and a random hexidecimal value. The value `a1d221c7fd97b08175b3872a37428e8c` in this example is a hash value calculated from the string `/home/cloudshell-user/job_attachments_devguide`, the directory where `script.sh` is located.

The manifest object contains the information for the input files on a specific root path uploaded to S3 as part of the job's submission. Download this manifest file (`aws s3 cp s3://$Q1_S3_BUCKET/<objectname>`). Its contents are similar to:

```json
{
    "hashAlg": "xxh128",
    "manifestVersion": "2023-03-03",
    "paths": [
        {
            "hash": "87cb19095dd5d78fcaf56384ef0e6241",
            "mtime": 1721147454416085,
            "path": "script.sh",
            "size": 39
        }
    ],
    "totalSize": 39
}
```

This indicates that the file `script.sh` was uploaded, and the hash of that file's contents is 87cb19095dd5d78fcaf56384ef0e6241. This hash value matches the value in the object name `DeadlineCloud/Data/87cb19095dd5d78fcaf56384ef0e6241.xxh128`. It is used by Deadline Cloud to know which object to download for this file's contents.

The full schema for this file is available in GitHub.

When you use the CreateJob operation you can set the location of the manifest objects. You can use the GetJob operation to see the location:

```json
{
    "attachments": {
        "file system": "COPIED",
        "manifests": [
            {
                "inputManifestHash": "5b0db3d311805ea8de7787b64cbbe8b3",
                "inputManifestPath": "<farm-id>/<queue-id>/Inputs/<guid>/
a1d221c7fd97b08175b3872a37428e8c_input",
                "rootPath": "/home/cloudshell-user/job_attachments_devguide",
                "rootPathFormat": "posix"
            }
        ]
    },
    ...
}
```

How Deadline Cloud chooses the files to upload

The files and directories that job attachments considers for upload to Amazon S3 as inputs to your job are:

- The values of all PATH-type job parameters defined in the job bundle's job template with a dataFlow value of IN or INOUT.

- The files and directories listed as inputs in the job bundle's asset references file.

If you submit a job with no storage profile, all of the files considered for uploading are uploaded. If you submit a job with a storage profile, files are not uploaded to Amazon S3 if they are located in the storage profile's SHARED-type file system locations that are also required file system locations for the queue. These locations are expected to be available on the worker hosts that run the job, so there is no need to upload them to S3.

In this example, you create SHARED file system locations in WSAll in your AWS CloudShell environment and then add files to those file system locations. Use the following command:

```
# Change the value of WSALL_ID to the identifier of the WSAll storage profile
WSALL_ID=sp-00112233445566778899aabbccddeeff

sudo mkdir -p /shared/common /shared/projects/project1 /shared/projects/project2
sudo chown -R cloudshell-user:cloudshell-user /shared

for d in /shared/common /shared/projects/project1 /shared/projects/project2; do
  echo "File contents for $d" > ${d}/file.txt
done
```

Next, add an asset references file to the job bundle that includes all the files that you created as inputs for the job. Use the following command:

```
cat > ${HOME}/job_attachments_devguide/asset_references.yaml << EOF
assetReferences:
  inputs:
    filenames:
    - /shared/common/file.txt
    directories:
    - /shared/projects/project1
    - /shared/projects/project2
EOF
```

Next, configure the Deadline Cloud CLI to submit jobs with the `WSAll` storage profile, and then submit the job bundle:

```
# Change the value of FARM_ID to your farm's identifier
FARM_ID=farm-00112233445566778899aabbccddeeff
# Change the value of QUEUE1_ID to queue Q1's identifier
QUEUE1_ID=queue-00112233445566778899aabbccddeeff
# Change the value of WSALL_ID to the identifier of the WSAll storage profile
WSALL_ID=sp-00112233445566778899aabbccddeeff

deadline config set settings.storage_profile_id $WSALL_ID

deadline bundle submit --farm-id $FARM_ID --queue-id $QUEUE1_ID
  job_attachments_devguide/
```

Deadline Cloud uploads two files to Amazon S3 when you submit the job. You can download the manifest objects for the job from S3 to see the uploaded files:

```
for manifest in $( \
  aws deadline get-job --farm-id $FARM_ID --queue-id $QUEUE1_ID --job-id $JOB_ID \
    --query 'attachments.manifests[].inputManifestPath' \
    | jq -r '.[]'
); do
  echo "Manifest object: $manifest"
  aws s3 cp --quiet s3://$Q1_S3_BUCKET/DeadlineCloud/Manifests/$manifest /dev/stdout |
 jq .
done
```

In this example, there is a single manifest file with the following contents:

```
{
    "hashAlg": "xxh128",
    "manifestVersion": "2023-03-03",
    "paths": [
        {
            "hash": "87cb19095dd5d78fcaf56384ef0e6241",
            "mtime": 1721147454416085,
            "path": "home/cloudshell-user/job_attachments_devguide/script.sh",
            "size": 39
        },
        {
```

```
            "hash": "af5a605a3a4e86ce7be7ac5237b51b79",
            "mtime": 1721163773582362,
            "path": "shared/projects/project2/file.txt",
            "size": 44
        }
    ],
    "totalSize": 83
}
```

Use the GetJob operation for the manifest to see that the rootPath is "/".

```
aws deadline get-job --farm-id $FARM_ID --queue-id $QUEUE1_ID --job-id $JOB_ID --query
 'attachments.manifests[*]'
```

The root path for set of input files is always the longest common subpath of those files. If your job
was submitted from Windows instead and there are input files with no common subpath because
they were on different drives, you see a separate root path on each drive. The paths in a manifest
are always relative to the root path of the manifest, so the input files that were uploaded are:

- /home/cloudshell-user/job_attachments_devguide/script.sh – The script file in the
  job bundle.

- /shared/projects/project2/file.txt – The file in a SHARED file system location in the
  WSAll storage profile that is **not** in the list of required file system locations for queue Q1.

The files in file system locations FSCommon (/shared/common/file.txt) and FS1 (/shared/
projects/project1/file.txt) are not in the list. This is because those file system locations
are SHARED in the WSAll storage profile and they both are in the list of required file system
locations in queue Q1.

You can see the file system locations considered SHARED for a job that is submitted with a
particular storage profile with the GetStorageProfileForQueue operation. To query for storage
profile WSAll for queue Q1 use the following command:

```
aws deadline get-storage-profile --farm-id $FARM_ID --storage-profile-id $WSALL_ID

aws deadline get-storage-profile-for-queue --farm-id $FARM_ID --queue-id $QUEUE1_ID --
storage-profile-id $WSALL_ID
```

How jobs find job attachment input files

For a job to use the files that Deadline Cloud uploads to Amazon S3 using job attachments, your job needs those files available through the file system on the worker hosts. When a [session](#) for your job runs on a worker host, Deadline Cloud downloads the input files for the job into a temporary directory on the worker host's local drive and adds path mapping rules for each of the job's root paths to its file system location on the local drive.

For this example, start the Deadline Cloud worker agent in an AWS CloudShell tab. Let any previously submitted jobs finish running, and then delete the job logs from the logs directory:

```
rm -rf ~/devdemo-logs/queue-*
```

The following script modifies the job bundle to show all files in the session's temporary working directory and the contents of the path mapping rules file, and then submits a job with the modified bundle:

```
# Change the value of FARM_ID to your farm's identifier
FARM_ID=farm-00112233445566778899aabbccddeeff
# Change the value of QUEUE1_ID to queue Q1's identifier
QUEUE1_ID=queue-00112233445566778899aabbccddeeff
# Change the value of WSALL_ID to the identifier of the WSAll storage profile
WSALL_ID=sp-00112233445566778899aabbccddeeff

deadline config set settings.storage_profile_id $WSALL_ID

cat > ~/job_attachments_devguide/script.sh << EOF
#!/bin/bash

echo "Session working directory is: \$(pwd)"
echo
echo "Contents:"
find . -type f
echo
echo "Path mapping rules file: \$1"
jq . \$1
EOF

cat > ~/job_attachments_devguide/template.yaml << EOF
specificationVersion: jobtemplate-2023-09
name: "Job Attachments Explorer"
parameterDefinitions:
```

```
- name: ScriptFile
  type: PATH
  default: script.sh
  dataFlow: IN
  objectType: FILE
steps:
- name: Step
  script:
    actions:
      onRun:
        command: /bin/bash
        args:
        - "{{Param.ScriptFile}}"
        - "{{Session.PathMappingRulesFile}}"
EOF

deadline bundle submit --farm-id $FARM_ID --queue-id $QUEUE1_ID
  job_attachments_devguide/
```

You can look at the log of the job's run after it has been run by the worker in your AWS CloudShell environment:

```
cat demoenv-logs/queue-*/session*.log
```

The log shows that the first thing that occurs in the session is the two input files for the job are downloaded to the worker:

```
2024-07-17 01:26:37,824 INFO ============================================
2024-07-17 01:26:37,825 INFO --------- Job Attachments Download for Job
2024-07-17 01:26:37,825 INFO ============================================
2024-07-17 01:26:37,825 INFO Syncing inputs using Job Attachments
2024-07-17 01:26:38,116 INFO Downloaded 142.0 B / 186.0 B of 2 files (Transfer rate:
 0.0 B/s)
2024-07-17 01:26:38,174 INFO Downloaded 186.0 B / 186.0 B of 2 files (Transfer rate:
 733.0 B/s)
2024-07-17 01:26:38,176 INFO Summary Statistics for file downloads:
Processed 2 files totaling 186.0 B.
Skipped re-processing 0 files totaling 0.0 B.
Total processing time of 0.09752 seconds at 1.91 KB/s.
```

Next is the output from `script.sh` run by the job:

- The input files uploaded when the job was submitted are located under a directory whose name begins with "assetroot" in the session's temporary directory.

- The input files' paths have been relocated relative to the "assetroot" directory instead of relative to the root path for the job's input manifest ("/").

- The path mapping rules file contains an additional rule that remaps "/" to the absolute path of the "assetroot" directory.

For example:

```
2024-07-17 01:26:38,264 INFO Output:
2024-07-17 01:26:38,267 INFO Session working directory is: /sessions/session-5b33f
2024-07-17 01:26:38,267 INFO
2024-07-17 01:26:38,267 INFO Contents:
2024-07-17 01:26:38,269 INFO ./tmp_xdhbsdo.sh
2024-07-17 01:26:38,269 INFO ./tmpdi00052b.json
2024-07-17 01:26:38,269 INFO ./assetroot-assetroot-3751a/shared/projects/project2/
file.txt
2024-07-17 01:26:38,269 INFO ./assetroot-assetroot-3751a/home/cloudshell-user/
job_attachments_devguide/script.sh
2024-07-17 01:26:38,269 INFO
2024-07-17 01:26:38,270 INFO Path mapping rules file: /sessions/session-5b33f/
tmpdi00052b.json
2024-07-17 01:26:38,282 INFO {
2024-07-17 01:26:38,282 INFO    "version": "pathmapping-1.0",
2024-07-17 01:26:38,282 INFO    "path_mapping_rules": [
2024-07-17 01:26:38,282 INFO      {
2024-07-17 01:26:38,282 INFO        "source_path_format": "POSIX",
2024-07-17 01:26:38,282 INFO        "source_path": "/shared/projects/project1",
2024-07-17 01:26:38,283 INFO        "destination_path": "/mnt/projects/project1"
2024-07-17 01:26:38,283 INFO      },
2024-07-17 01:26:38,283 INFO      {
2024-07-17 01:26:38,283 INFO        "source_path_format": "POSIX",
2024-07-17 01:26:38,283 INFO        "source_path": "/shared/common",
2024-07-17 01:26:38,283 INFO        "destination_path": "/mnt/common"
2024-07-17 01:26:38,283 INFO      },
2024-07-17 01:26:38,283 INFO      {
2024-07-17 01:26:38,283 INFO        "source_path_format": "POSIX",
2024-07-17 01:26:38,283 INFO        "source_path": "/",
2024-07-17 01:26:38,283 INFO        "destination_path": "/sessions/session-5b33f/
assetroot-assetroot-3751a"
2024-07-17 01:26:38,283 INFO      }
```

```
2024-07-17 01:26:38,283 INFO   ]
2024-07-17 01:26:38,283 INFO }
```

> **ⓘ Note**
>
> If the job you submit has multiple manifests with different root paths, there is a different "assetroot"-named directory for each of the root paths.

If you need to reference the relocated file system location of one of your input files, directories, or file system locations you can either process the path mapping rules file in your job and perform the remapping yourself, or add a PATH type job parameter to the job template in your job bundle and pass the value that you need to remap as the value of that parameter. For example, the following example modifies the job bundle to have one of these job parameters and then submits a job with the file system location /shared/projects/project2 as its value:

```
cat > ~/job_attachments_devguide/template.yaml << EOF
specificationVersion: jobtemplate-2023-09
name: "Job Attachments Explorer"
parameterDefinitions:
- name: LocationToRemap
  type: PATH
steps:
- name: Step
  script:
    actions:
      onRun:
        command: /bin/echo
        args:
        - "The location of {{RawParam.LocationToRemap}} in the session is
 {{Param.LocationToRemap}}"
EOF

deadline bundle submit --farm-id $FARM_ID --queue-id $QUEUE1_ID
 job_attachments_devguide/ \
   -p LocationToRemap=/shared/projects/project2
```

The log file for this job's run contains its output:

```
2024-07-17 01:40:35,283 INFO Output:
```

```
2024-07-17 01:40:35,284 INFO The location of /shared/projects/project2 in the session
  is /sessions/session-5b33f/assetroot-assetroot-3751a
```

Getting output files from a job

This example shows how Deadline Cloud identifies the output files that your jobs generate, decides whether to upload those files to Amazon S3, and how you can get those output files on your workstation.

Use the `job_attachments_devguide_output` job bundle instead of the `job_attachments_devguide` job bundle for this example. Start by making a copy of the bundle in your AWS CloudShell environment from your clone of the Deadline Cloud samples GitHub repository:

```
cp -r deadline-cloud-samples/job_bundles/job_attachments_devguide_output ~/
```

The important difference between this job bundle and the `job_attachments_devguide` job bundle is the addition of a new job parameter in the job template:

```
...
parameterDefinitions:
...
- name: OutputDir
  type: PATH
  objectType: DIRECTORY
  dataFlow: OUT
  default: ./output_dir
  description: This directory contains the output for all steps.
...
```

The `dataFlow` property of the parameter has the value `OUT`. Deadline Cloud uses the value of `dataFlow` job parameters with a value of `OUT` or `INOUT` as outputs of your job. If the file system location passed as a value to these kinds of job parameters is remapped to a local file system location on the worker that runs the job, then Deadline Cloud will look for new files at the location and upload those to Amazon S3 as job outputs.

To see how this works, first start the Deadline Cloud worker agent in an AWS CloudShell tab. Let any previously submitted jobs finish running. Then delete the job logs from the logs directory:

```
rm -rf ~/devdemo-logs/queue-*
```

Next, submit a job with this job bundle. After the worker running in your CloudShell runs, look at the logs:

```
# Change the value of FARM_ID to your farm's identifier
FARM_ID=farm-00112233445566778899aabbccddeeff
# Change the value of QUEUE1_ID to queue Q1's identifier
QUEUE1_ID=queue-00112233445566778899aabbccddeeff
# Change the value of WSALL_ID to the identifier of the WSAll storage profile
WSALL_ID=sp-00112233445566778899aabbccddeeff

deadline config set settings.storage_profile_id $WSALL_ID

deadline bundle submit --farm-id $FARM_ID --queue-id $QUEUE1_ID ./
job_attachments_devguide_output
```

The log shows that a file was detected as output and uploaded to Amazon S3:

```
2024-07-17 02:13:10,873 INFO ---------------------------------------------
2024-07-17 02:13:10,873 INFO Uploading output files to Job Attachments
2024-07-17 02:13:10,873 INFO ---------------------------------------------
2024-07-17 02:13:10,873 INFO Started syncing outputs using Job Attachments
2024-07-17 02:13:10,955 INFO Found 1 file totaling 117.0 B in output directory: /
sessions/session-7efa/assetroot-assetroot-3751a/output_dir
2024-07-17 02:13:10,956 INFO Uploading output manifest to
 DeadlineCloud/Manifests/farm-0011/queue-2233/job-4455/step-6677/
task-6677-0/2024-07-17T02:13:10.835545Z_sessionaction-8899-1/
c6808439dfc59f86763aff5b07b9a76c_output
2024-07-17 02:13:10,988 INFO Uploading 1 output file to S3: s3BucketName/DeadlineCloud/
Data
2024-07-17 02:13:11,011 INFO Uploaded 117.0 B / 117.0 B of 1 file (Transfer rate: 0.0
 B/s)
2024-07-17 02:13:11,011 INFO Summary Statistics for file uploads:
Processed 1 file totaling 117.0 B.
Skipped re-processing 0 files totaling 0.0 B.
Total processing time of 0.02281 seconds at 5.13 KB/s.
```

The log also shows that Deadline Cloud created a new manifest object in the Amazon S3 bucket configured for use by job attachments on queue Q1. The name of the manifest object is derived from the farm, queue, job, step, task, timestamp, and `sessionaction` identifiers of the task that generated the output. Download this manifest file to see where Deadline Cloud placed the output files for this task:

```
# The name of queue `Q1`'s job attachments S3 bucket
Q1_S3_BUCKET=$(
  aws deadline get-queue --farm-id $FARM_ID --queue-id $QUEUE1_ID \
    --query 'jobAttachmentSettings.s3BucketName' | tr -d '"'
)

# Fill this in with the object name from your log
OBJECT_KEY="DeadlineCloud/Manifests/..."

aws s3 cp --quiet s3://$Q1_S3_BUCKET/$OBJECT_KEY /dev/stdout | jq .
```

The manifest looks like:

```
{
  "hashAlg": "xxh128",
  "manifestVersion": "2023-03-03",
  "paths": [
    {
      "hash": "34178940e1ef9956db8ea7f7c97ed842",
      "mtime": 1721182390859777,
      "path": "output_dir/output.txt",
      "size": 117
    }
  ],
  "totalSize": 117
}
```

This shows that the content of the output file is saved to Amazon S3 the same way that job input files are saved. Similar to input files, the output file is stored in S3 with an object name containing the hash of the file and the prefix `DeadlineCloud/Data`.

```
$ aws s3 ls --recursive s3://$Q1_S3_BUCKET | grep 34178940e1ef9956db8ea7f7c97ed842
2024-07-17 02:13:11        117 DeadlineCloud/
Data/34178940e1ef9956db8ea7f7c97ed842.xxh128
```

You can download the output of a job to your workstation using the Deadline Cloud monitor or the Deadline Cloud CLI:

```
deadline job download-output --farm-id $FARM_ID --queue-id $QUEUE1_ID --job-id $JOB_ID
```

The value of the `OutputDir` job parameter in the submitted job is `./output_dir`, so the output
are downloaded to a directory called `output_dir` within the job bundle directory. If you specified
an absolute path or different relative location as the value for `OutputDir`, then the output files
would be downloaded to that location instead.

```
$ deadline job download-output --farm-id $FARM_ID --queue-id $QUEUE1_ID --job-id
 $JOB_ID
Downloading output from Job 'Job Attachments Explorer: Output'

Summary of files to download:
    /home/cloudshell-user/job_attachments_devguide_output/output_dir/output.txt (1
 file)

You are about to download files which may come from multiple root directories. Here are
 a list of the current root directories:
[0] /home/cloudshell-user/job_attachments_devguide_output
> Please enter the index of root directory to edit, y to proceed without changes, or n
 to cancel the download (0, y, n) [y]:

Downloading Outputs  [################################]  100%
Download Summary:
    Downloaded 1 files totaling 117.0 B.
    Total download time of 0.14189 seconds at 824.0 B/s.
    Download locations (total file counts):
        /home/cloudshell-user/job_attachments_devguide_output (1 file)
```

Using files from a step in a dependent step

This example shows how one step in a job can access the outputs from a step that it depends on in
the same job.

To make the outputs of one step available to another, Deadline Cloud adds additional actions to a
session to download those outputs before running tasks in the session. You tell it which steps to
download the outputs from by declaring those steps as dependencies of the step that needs to use
the outputs.

Use the `job_attachments_devguide_output` job bundle for this example. Start by making a
copy in your AWS CloudShell environment from your clone of the Deadline Cloud samples GitHub
repository. Modify it to add a dependent step that only runs after the existing step and uses that
step's output:

```
cp -r deadline-cloud-samples/job_bundles/job_attachments_devguide_output ~/

cat >> job_attachments_devguide_output/template.yaml << EOF
- name: DependentStep
  dependencies:
  - dependsOn: Step
  script:
    actions:
      onRun:
        command: /bin/cat
        args:
        - "{{Param.OutputDir}}/output.txt"
EOF
```

The job created with this modified job bundle runs as two separate sessions, one for the task in the step "Step" and then a second for the task in the step "DependentStep".

First start the Deadline Cloud worker agent in an CloudShell tab. Let any previously submitted jobs finish running, then delete the job logs from the logs directory:

```
rm -rf ~/devdemo-logs/queue-*
```

Next, submit a job using the modified `job_attachments_devguide_output` job bundle. Wait for it to finish running on the worker in your CloudShell environment. Look at the logs for the two sessions:

```
# Change the value of FARM_ID to your farm's identifier
FARM_ID=farm-00112233445566778899aabbccddeeff
# Change the value of QUEUE1_ID to queue Q1's identifier
QUEUE1_ID=queue-00112233445566778899aabbccddeeff
# Change the value of WSALL_ID to the identifier of the WSAll storage profile
WSALL_ID=sp-00112233445566778899aabbccddeeff

deadline config set settings.storage_profile_id $WSALL_ID

deadline bundle submit --farm-id $FARM_ID --queue-id $QUEUE1_ID ./
job_attachments_devguide_output

# Wait for the job to finish running, and then:

cat demoenv-logs/queue-*/session-*
```

In the session log for the task in the step named `DependentStep`, there are two separate download actions run:

```
2024-07-17 02:52:05,666 INFO ============================================
2024-07-17 02:52:05,666 INFO --------- Job Attachments Download for Job
2024-07-17 02:52:05,667 INFO ============================================
2024-07-17 02:52:05,667 INFO Syncing inputs using Job Attachments
2024-07-17 02:52:05,928 INFO Downloaded 207.0 B / 207.0 B of 1 file (Transfer rate: 0.0
 B/s)
2024-07-17 02:52:05,929 INFO Summary Statistics for file downloads:
Processed 1 file totaling 207.0 B.
Skipped re-processing 0 files totaling 0.0 B.
Total processing time of 0.03954 seconds at 5.23 KB/s.

2024-07-17 02:52:05,979 INFO
2024-07-17 02:52:05,979 INFO ============================================
2024-07-17 02:52:05,979 INFO --------- Job Attachments Download for Step
2024-07-17 02:52:05,979 INFO ============================================
2024-07-17 02:52:05,980 INFO Syncing inputs using Job Attachments
2024-07-17 02:52:06,133 INFO Downloaded 117.0 B / 117.0 B of 1 file (Transfer rate: 0.0
 B/s)
2024-07-17 02:52:06,134 INFO Summary Statistics for file downloads:
Processed 1 file totaling 117.0 B.
Skipped re-processing 0 files totaling 0.0 B.
Total processing time of 0.03227 seconds at 3.62 KB/s.
```

The first action downloads the `script.sh` file used by the step named "Step." The second action downloads the outputs from that step. Deadline Cloud determines which files to download by using the output manifest generated by that step as an input manifest.

Late in the same log, you can see the output from the step named "DependentStep":

```
2024-07-17 02:52:06,213 INFO Output:
2024-07-17 02:52:06,216 INFO Script location: /sessions/session-5b33f/
assetroot-assetroot-3751a/script.sh
```

Security in Deadline Cloud

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS 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 [AWS Compliance Programs](#). To learn about the compliance programs that apply to AWS Deadline Cloud, see [AWS services in Scope by Compliance Program](#).

- **Security in the cloud** – Your responsibility is determined by the AWS 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 Deadline Cloud. The following topics show you how to configure Deadline Cloud to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Deadline Cloud resources.

Topics

- [Data protection in Deadline Cloud](#)
- [Identity and Access Management in Deadline Cloud](#)
- [Compliance validation for Deadline Cloud](#)
- [Resilience in Deadline Cloud](#)
- [Infrastructure security in Deadline Cloud](#)
- [Configuration and vulnerability analysis in Deadline Cloud](#)
- [Cross-service confused deputy prevention](#)
- [Access AWS Deadline Cloud using an interface endpoint (AWS PrivateLink)](#)
- [Security best practices for Deadline Cloud](#)

Data protection in Deadline Cloud

The AWS [shared responsibility model](#) applies to data protection in AWS Deadline Cloud. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS 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 AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#). For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS 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 AWS resources. We require TLS 1.2 and recommend TLS 1.3.

- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see [Working with CloudTrail trails](#) in the *AWS CloudTrail User Guide*.

- Use AWS encryption solutions, along with all default security controls within AWS 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 140-3 validated cryptographic modules when accessing AWS 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) 140-3](#).

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 Deadline Cloud or other AWS services using the console, API, AWS CLI, or AWS 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.

The data entered into name fields in Deadline Cloud job templates may also be included in billing or diagnostic logs and should not contain confidential or sensitive information.

Topics

- [Encryption at rest](#)

- [Encryption in transit](#)

- [Key management](#)

- [Inter-network traffic privacy](#)

- [Opt out](#)

Encryption at rest

AWS Deadline Cloud protects sensitive data by encrypting it at rest using encryption keys stored in [AWS Key Management Service (AWS KMS)](#). Encryption at rest is available in all AWS Regions where Deadline Cloud is available.

Encrypting data means sensitive data saved on disks isn't readable by a user or application without a valid key. Only a party with a valid managed key can decrypt the data.

For information about how Deadline Cloud uses AWS KMS for encrypting data at rest, see [Key management](#).

Encryption in transit

For data in transit, AWS Deadline Cloud uses Transport Layer Security (TLS) 1.2 or 1.3 to encrypt data sent between the service and workers. We require TLS 1.2 and recommend TLS 1.3. Additionally, if you use a virtual private cloud (VPC), you can use AWS PrivateLink to establish a private connection between your VPC and Deadline Cloud.

Key management

When creating a new farm, you can choose one of the following keys to encrypt your farm data:

- **AWS owned KMS key** – Default encryption type if you don't specify a key when you create the farm. The KMS key is owned by AWS Deadline Cloud. You can't view, manage, or use AWS owned keys. However, you don't need to take any action to protect the keys that encrypt your data. For more information, see [AWS owned keys](#) in the *AWS Key Management Service developer guide*.

- **Customer managed KMS key** – You specify a customer managed key when you create a farm. All of the content within the farm is encrypted with the KMS key. The key is stored in your account

and is created, owned, and managed by you and AWS KMS charges apply. You have full control over the KMS key. You can perform such tasks as:

- Establishing and maintaining key polices

- Establishing and maintaining IAM policies and grants

- Enabling and disabling key policies

- Adding tags

- Creating key aliases

You can't manually rotate a customer owned key used with a Deadline Cloud farm. Automatic rotation of the key is supported.

For more information, see Customer owned keys in the *AWS Key Management Service Developer Guide*.

To create a customer managed key, follow the steps for Creating symmetric customer managed keys in the *AWS Key Management Service Developer Guide*.

How Deadline Cloud use AWS KMS grants

Deadline Cloud requires a grant to use your customer managed key. When you create a farm encrypted with a customer managed key, Deadline Cloud creates a grant on your behalf by sending a `CreateGrant` request to AWS KMS to get access to the KMS key that you specified.

Deadline Cloud uses multiple grants. Each grant is used by a different part of Deadline Cloud that needs to encrypt or decrypt your data. Deadline Cloud also uses grants to allow access to other AWS services used to store data on your behalf, such as Amazon Simple Storage Service, Amazon Elastic Block Store, or OpenSearch.

Grants that enable Deadline Cloud to manage machines in a service-managed fleet include a Deadline Cloud account number and role in the `GranteePrincipal` instead of a service principal. While not typical, this is necessary to encrypt Amazon EBS volumes for workers in service-managed fleets using the customer managed KMS key specified for the farm.

Customer managed key policy

Key policies control access to your customer managed key. Each key must have exactly one key policy that contains statements that determine who can use the key and how they can use it. When

you create you customer managed key, you can specify a key policy. For more information, see [Managing access to customer managed keys](#) in the *AWS Key Management Service Developer Guide*.

Minimal IAM policy for CreateFarm

To use your customer managed key to create farms using the console or the `CreateFarm` API operation, the following AWS KMS API operations must be permitted:

- `kms:CreateGrant` – Adds a grant to a customer managed key. Grants console access to a specified AWS KMS key. For more informations, see [Using grants](#) in the *AWS Key Management Service developer guide*.

- `kms:Decrypt` – Allows Deadline Cloud to decrypt data in the farm.

- `kms:DescribeKey` – Provides the customer managed key details to allow Deadline Cloud to validate the key.

- `kms:GenerateDataKey` – Allows Deadline Cloud to encrypt data using a unique data key.

The following policy statement grants the necessary permissions for the `CreateFarm` operation.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "DeadlineCreateGrants",
            "Effect": "Allow",
            "Action": [
                "kms:Decrypt",
                "kms:GenerateDataKey",
                "kms:CreateGrant",
                "kms:DescribeKey"
            ],
            "Resource": "arn:aws::kms:us-west-2:111122223333:key/1234567890abcdef0",
            "Condition": {
                "StringEquals": {
                    "kms:ViaService": "deadline.us-west-2.amazonaws.com"
                }
            }
        }
    ]
}
```

Minimal IAM policy for read-only operations

To use your customer managed key for read-only Deadline Cloud operations, such getting information about farms, queues, and fleets. The following AWS KMS API operations must be permitted:

- kms:Decrypt – Allows Deadline Cloud to decrypt data in the farm.
- kms:DescribeKey – Provides the customer managed key details to allow Deadline Cloud to validate the key.

The following policy statement grants the necessary permissions for read-only operations.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "DeadlineReadOnly",
            "Effect": "Allow",
            "Action": [
                "kms:Decrypt",
                "kms:DescribeKey"
            ],
            "Resource": "arn:aws::kms:us-west-2:111122223333:key/a1b2c3d4-5678-90ab-
cdef-EXAMPLE11111",
            "Condition": {
                "StringEquals": {
                    "kms:ViaService": "deadline.us-west-2.amazonaws.com"
                }
            }
        }
    ]
}
```

Minimal IAM policy for read-write operations

To use your customer managed key for read-write Deadline Cloud operations, such as creating and updating farms, queues, and fleets. The following AWS KMS API operations must be permitted:

- kms:Decrypt – Allows Deadline Cloud to decrypt data in the farm.
- kms:DescribeKey – Provides the customer managed key details to allow Deadline Cloud to validate the key.

- `kms:GenerateDataKey` – Allows Deadline Cloud to encrypt data using a unique data key.

The following policy statement grants the necessary permissions for the `CreateFarm` operation.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "DeadlineReadWrite",
            "Effect": "Allow",
            "Action": [
                "kms:Decrypt",
                "kms:DescribeKey",
                "kms:GenerateDataKey",
            ],
            "Resource": "arn:aws::kms:us-west-2:111122223333:key/a1b2c3d4-5678-90ab-cdef-EXAMPLE11111",
            "Condition": {
                "StringEquals": {
                    "kms:ViaService": "deadline.us-west-2.amazonaws.com"
                }
            }
        }
    ]
}
```

Monitoring your encryption keys

When you use an AWS KMS customer managed key with your Deadline Cloud farms, you can use
AWS CloudTrail or Amazon CloudWatch Logs to track requests that Deadline Cloud sends to AWS
KMS.

CloudTrail event for grants

The following example CloudTrail event occurs when grants are created, typically when you call the
`CreateFarm`, `CreateMonitor`, or `CreateFleet` operation.

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "AROAIGDTESTANDEXAMPLE:SampleUser01",
```

```
        "arn": "arn:aws::sts::111122223333:assumed-role/Admin/SampleUser01",
        "accountId": "111122223333",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE3",
        "sessionContext": {
            "sessionIssuer": {
                "type": "Role",
                "principalId": "AROAIGDTESTANDEXAMPLE",
                "arn": "arn:aws::iam::111122223333:role/Admin",
                "accountId": "111122223333",
                "userName": "Admin"
            },
            "webIdFederationData": {},
            "attributes": {
                "creationDate": "2024-04-23T02:05:26Z",
                "mfaAuthenticated": "false"
            }
        },
        "invokedBy": "deadline.amazonaws.com"
    },
    "eventTime": "2024-04-23T02:05:35Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "CreateGrant",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "deadline.amazonaws.com",
    "userAgent": "deadline.amazonaws.com",
    "requestParameters": {
        "operations": [
            "CreateGrant",
            "Decrypt",
            "DescribeKey",
            "Encrypt",
            "GenerateDataKey"
        ],
        "constraints": {
            "encryptionContextSubset": {
                "aws:deadline:farmId": "farm-abcdef12345678900987654321fedcba",
                "aws:deadline:accountId": "111122223333"
            }
        },
        "granteePrincipal": "deadline.amazonaws.com",
        "keyId": "arn:aws::kms:us-west-2:111122223333:key/a1b2c3d4-5678-90ab-cdef-
EXAMPLE11111",
        "retiringPrincipal": "deadline.amazonaws.com"
    },
```

```
    "responseElements": {
        "grantId": "6bbe819394822a400fe5e3a75d0e9ef16c1733143fff0c1fc00dc7ac282a18a0",
        "keyId": "arn:aws::kms:us-west-2:111122223333:key/a1b2c3d4-5678-90ab-cdef-
EXAMPLE11111"
    },
    "requestID": "a1b2c3d4-5678-90ab-cdef-EXAMPLE22222",
    "eventID": "a1b2c3d4-5678-90ab-cdef-EXAMPLE33333",
    "readOnly": false,
    "resources": [
        {
            "accountId": "AWS Internal",
            "type": "AWS::KMS::Key",
            "ARN": "arn:aws::kms:us-west-2:111122223333:key/a1b2c3d4-5678-90ab-cdef-
EXAMPLE44444"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "111122223333",
    "eventCategory": "Management"
}
```

CloudTrail event for decryption

The following example CloudTrail event occurs when decrypting values using the customer managed KMS key.

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "AROAIGDTESTANDEXAMPLE:SampleUser01",
        "arn": "arn:aws::sts::111122223333:assumed-role/SampleRole/SampleUser01",
        "accountId": "111122223333",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "sessionContext": {
            "sessionIssuer": {
                "type": "Role",
                "principalId": "AROAIGDTESTANDEXAMPLE",
                "arn": "arn:aws::iam::111122223333:role/SampleRole",
                "accountId": "111122223333",
                "userName": "SampleRole"
            },
```

```
            "webIdFederationData": {},
            "attributes": {
                "creationDate": "2024-04-23T18:46:51Z",
                "mfaAuthenticated": "false"
            }
        },
        "invokedBy": "deadline.amazonaws.com"
    },
    "eventTime": "2024-04-23T18:51:44Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "Decrypt",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "deadline.amazonaws.com",
    "userAgent": "deadline.amazonaws.com",
    "requestParameters": {
        "encryptionContext": {
            "aws:deadline:farmId": "farm-abcdef12345678900987654321fedcba",
            "aws:deadline:accountId": "111122223333",
            "aws-crypto-public-key": "AotL+SAMPLEVALUEiOMEXAMPLEaaqNOTREALaGTESTONLY
+p/5H+EuKd4Q=="
        },
        "encryptionAlgorithm": "SYMMETRIC_DEFAULT",
        "keyId": "arn:aws::kms:us-west-2:111122223333:key/a1b2c3d4-5678-90ab-cdef-
EXAMPLE11111"
    },
    "responseElements": null,
    "requestID": "aaaaaaaa-bbbb-cccc-dddd-eeeeeeffffff",
    "eventID": "ffffffff-eeee-dddd-cccc-bbbbbbaaaaaa",
    "readOnly": true,
    "resources": [
        {
            "accountId": "111122223333",
            "type": "AWS::KMS::Key",
            "ARN": "arn:aws::kms:us-west-2:111122223333:key/a1b2c3d4-5678-90ab-cdef-
EXAMPLE11111"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "111122223333",
    "eventCategory": "Management"
}
```

CloudTrail event for encryption

The following example CloudTrail event occurs when encrypting values using the customer managed KMS key.

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "AROAIGDTESTANDEXAMPLE:SampleUser01",
        "arn": "arn:aws::sts::111122223333:assumed-role/SampleRole/SampleUser01",
        "accountId": "111122223333",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "sessionContext": {
            "sessionIssuer": {
                "type": "Role",
                "principalId": "AROAIGDTESTANDEXAMPLE",
                "arn": "arn:aws::iam::111122223333:role/SampleRole",
                "accountId": "111122223333",
                "userName": "SampleRole"
            },
            "webIdFederationData": {},
            "attributes": {
                "creationDate": "2024-04-23T18:46:51Z",
                "mfaAuthenticated": "false"
            }
        },
        "invokedBy": "deadline.amazonaws.com"
    },
    "eventTime": "2024-04-23T18:52:40Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "GenerateDataKey",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "deadline.amazonaws.com",
    "userAgent": "deadline.amazonaws.com",
    "requestParameters": {
        "numberOfBytes": 32,
        "encryptionContext": {
            "aws:deadline:farmId": "farm-abcdef12345678900987654321fedcba",
            "aws:deadline:accountId": "111122223333",
            "aws-crypto-public-key": "AotL+SAMPLEVALUEiOMEXAMPLEaaqNOTREALaGTESTONLY
+p/5H+EuKd4Q=="
        },
```

```
        "keyId": "arn:aws::kms:us-
west-2:111122223333:key/abcdef12-3456-7890-0987-654321fedcba"
    },
    "responseElements": null,
    "requestID": "a1b2c3d4-5678-90ab-cdef-EXAMPLE11111",
    "eventID": "a1b2c3d4-5678-90ab-cdef-EXAMPLE22222",
    "readOnly": true,
    "resources": [
        {
            "accountId": "111122223333",
            "type": "AWS::KMS::Key",
            "ARN": "arn:aws::kms:us-west-2:111122223333:key/a1b2c3d4-5678-90ab-cdef-
EXAMPLE33333"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "111122223333",
    "eventCategory": "Management"
}
```

Deleting a customer managed KMS key

Deleting a customer managed KMS key in AWS Key Management Service (AWS KMS) is destructive and potentially dangerous. It irreversibly deletes the key material and all metadata associated with the key. After a customer managed KMS key is deleted, you can no longer decrypt the data that was encrypted by that key. This means that the data becomes unrecoverable.

This is why AWS KMS gives customers a waiting period of up to 30 days before deleting the KMS key. The default waiting period is 30 days.

About the waiting period

Because it's destructive and potentially dangerous to delete a customer managed KMS key, we require that you set a waiting period of 7–30 days. The default waiting period is 30 days.

However, the actual waiting period might be up to 24 hours longer than the period you scheduled. To get the actual date and time when the key will be deleted, use the DescribeKey operation. You can also see the scheduled deletion date of a key in the AWS KMS console on the key's detail page, in the **General configuration** section. Notice the time zone.

During the waiting period, the customer managed key's status and key state is **Pending deletion**.

- A customer managed KMS key that is pending deletion can't be used in any [cryptographic operations](#).

- AWS KMS doesn't [rotate the backing keys](#) of customer managed KMS keys that are pending deletion.

For more information about deleting a customer managed KMS key, see [Deleting customer master keys](#) in the *AWS Key Management Service Developer Guide*.

Inter-network traffic privacy

AWS Deadline Cloud supports Amazon Virtual Private Cloud (Amazon VPC) to secure connections. Amazon VPC provides features that you can use to increase and monitor the security for your virtual private cloud (VPC).

You can set up a customer-managed fleet (CMF) with Amazon Elastic Compute Cloud (Amazon EC2) instances that run inside a VPC. By deploying Amazon VPC endpoints to use AWS PrivateLink, traffic between workers in your CMF and the Deadline Cloud endpoint stays within your VPC. Furthermore, you can configure your VPC to restrict internet access to your instances.

In service-managed fleets, workers aren't reachable from the internet, but they do have internet access and connect to the Deadline Cloud service over the internet.

Opt out

AWS Deadline Cloud collects certain operational information to help us develop and improve Deadline Cloud. The collected data includes things such as your AWS account ID and user ID, so that we can correctly identify you if you have an issue with the Deadline Cloud. We also collect Deadline Cloud specific information, such as Resource IDs (a FarmID or QueueID when applicable), the product name (for example, JobAttachments, WorkerAgent, and more) and the product version.

You can choose to opt out from this data collection using application configuration. Each computer interacting with Deadline Cloud, both client workstations and fleet workers, needs to opt out separately.

Deadline Cloud monitor - desktop

Deadline Cloud monitor - desktop collects operational information, such as when crashes occur and when the application is opened, to help us know when you are having problems with the

application. To opt out from the collection of this operational information, go to the settings page and clear **Turn on data collection to measure Deadline Cloud Monitor's performance**.

After you opt out, the desktop monitor no longer sends the operational data. Any previously collected data is retained and may still be used to improve the service. For more information, see Data Privacy FAQ.

AWS Deadline Cloud CLI and Tools

The AWS Deadline Cloud CLI, submitters, and worker agent all collect operational information such as when crashes occur and when jobs are submitted to help us know when you are having problems with these applications. To opt out from the collection of this operational information, use any of the following methods:

- In the terminal, enter **`deadline config set telemetry.opt_out true`**.

  This will opt out the CLI, submitters, and worker agent when running as the current user.

- When installing the Deadline Cloud worker agent, add the **`--telemetry-opt-out`** command line argument. For example, **`./install.sh --farm-id $FARM_ID --fleet-id $FLEET_ID --telemetry-opt-out`**.

- Before running the worker agent, CLI, or submitter, set an environment variable: **`DEADLINE_CLOUD_TELEMETRY_OPT_OUT=true`**

After you opt out, the Deadline Cloud tools no longer send the operational data. Any previously collected data is retained and may still be used to improve the service. For more information, see Data Privacy FAQ.

Identity and Access Management in Deadline Cloud

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Deadline Cloud resources. IAM is an AWS service that you can use with no additional charge.

Topics

- Audience

- [Authenticating with identities](#)

- [Managing access using policies](#)

- [How Deadline Cloud works with IAM](#)

- [Identity-based policy examples for Deadline Cloud](#)

- [AWS managed policies for Deadline Cloud](#)

- [Troubleshooting AWS Deadline Cloud identity and access](#)

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in Deadline Cloud.

Service user – If you use the Deadline Cloud service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Deadline Cloud 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 Deadline Cloud, see [Troubleshooting AWS Deadline Cloud identity and access](#).

Service administrator – If you're in charge of Deadline Cloud resources at your company, you probably have full access to Deadline Cloud. It's your job to determine which Deadline Cloud 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 Deadline Cloud, see [How Deadline Cloud 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 Deadline Cloud. To view example Deadline Cloud identity-based policies that you can use in IAM, see [Identity-based policy examples for Deadline Cloud](#).

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be *authenticated* (signed in to AWS) as the AWS account root user, as an IAM user, or by assuming an IAM role.

You can sign in to AWS as a federated identity by using credentials provided through an identity source. AWS IAM Identity Center (IAM Identity Center) users, your company's single sign-on

authentication, and your Google or Facebook credentials are examples of federated identities. When you sign in as a federated identity, your administrator previously set up identity federation using IAM roles. When you access AWS by using federation, you are indirectly assuming a role.

Depending on the type of user you are, you can sign in to the AWS Management Console or the AWS access portal. For more information about signing in to AWS, see How to sign in to your AWS account in the *AWS Sign-In User Guide*.

If you access AWS programmatically, AWS 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 AWS tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see Signing AWS 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, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see Multi-factor authentication in the *AWS IAM Identity Center User Guide* and Using multi-factor authentication (MFA) in AWS in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS 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*.

Federated identity

As a best practice, require human users, including users that require administrator access, to use federation with an identity provider to access AWS services by using temporary credentials.

A *federated identity* is a user from your enterprise user directory, a web identity provider, the AWS Directory Service, the Identity Center directory, or any user that accesses AWS services by using credentials provided through an identity source. When federated identities access AWS accounts, they assume roles, and the roles provide temporary credentials.

For centralized access management, we recommend that you use AWS IAM Identity Center. You can create users and groups in IAM Identity Center, or you can connect and synchronize to a set of users and groups in your own identity source for use across all your AWS accounts and applications. For information about IAM Identity Center, see What is IAM Identity Center? in the *AWS IAM Identity Center User Guide*.

IAM users and groups

An *IAM user* is an identity within your AWS 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 AWS 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 AWS Management Console by switching roles. You can assume a role by calling an AWS CLI or AWS API operation or by using a custom URL. For more information about methods for using roles, see Methods to assume a role 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*. If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the permission set to a role in IAM. For information about permissions sets, see [Permission sets](#) in the *AWS IAM Identity Center 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 AWS 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 [Cross account resource access in IAM](#) in the *IAM User Guide*.

- **Cross-service access** – Some AWS services use features in other AWS 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 AWS, 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 AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS 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 AWS service](#) in the *IAM User Guide*.

  - **Service-linked role** – A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS 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 AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS 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 AWS by creating policies and attaching them to AWS identities or resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS 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 AWS 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 AWS 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 AWS Management Console, the AWS CLI, or the AWS 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 AWS account. Managed

policies include AWS 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 AWS services.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see Access control list (ACL) overview in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS 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 AWS Organizations. AWS Organizations is a service for grouping and centrally managing multiple AWS 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 AWS account root user. For more information about Organizations and SCPs, see [Service control policies](#) in the *AWS 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 AWS determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

How Deadline Cloud works with IAM

Before you use IAM to manage access to Deadline Cloud, learn what IAM features are available to use with Deadline Cloud.

IAM features you can use with AWS Deadline Cloud

IAM feature	Deadline Cloud support
[Identity-based policies](#)	Yes
[Resource-based policies](#)	No
[Policy actions](#)	Yes
[Policy resources](#)	Yes
[Policy condition keys (service-specific)](#)	Yes

IAM feature	Deadline Cloud support
ACLs	No
ABAC (tags in policies)	Yes
Temporary credentials	Yes
Forward access sessions (FAS)	Yes
Service roles	Yes
Service-linked roles	No

To get a high-level view of how Deadline Cloud and other AWS services work with most IAM features, see AWS services that work with IAM in the *IAM User Guide*.

Identity-based policies for Deadline Cloud

Supports identity-based policies: Yes

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*.

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. You can't specify the principal in an identity-based policy because it applies to the user or role to which it is attached. To learn about all of the elements that you can use in a JSON policy, see IAM JSON policy elements reference in the *IAM User Guide*.

Identity-based policy examples for Deadline Cloud

To view examples of Deadline Cloud identity-based policies, see Identity-based policy examples for Deadline Cloud.

Resource-based policies within Deadline Cloud

Supports resource-based policies: No

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 AWS services.

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 AWS accounts, an IAM administrator in the trusted account must also grant the principal entity (user or role) permission to access the resource. They 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 [Cross account resource access in IAM](#) in the *IAM User Guide*.

Policy actions for Deadline Cloud

Supports policy actions: Yes

Administrators can use AWS 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 AWS 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.

To see a list of Deadline Cloud actions, see [Actions defined by AWS Deadline Cloud](#) in the *Service Authorization Reference*.

Policy actions in Deadline Cloud use the following prefix before the action:

```
deadline
```

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [
      "deadline:action1",
      "deadline:action2"
        ]
```

To view examples of Deadline Cloud identity-based policies, see Identity-based policy examples for Deadline Cloud.

Policy resources for Deadline Cloud

Supports policy resources: Yes

Administrators can use AWS 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": "*"
```

To see a list of Deadline Cloud resource types and their ARNs, see Resources defined by AWS Deadline Cloud in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see Actions defined by AWS Deadline Cloud.

To view examples of Deadline Cloud identity-based policies, see Identity-based policy examples for Deadline Cloud.

Policy condition keys for Deadline Cloud

Supports service-specific policy condition keys: Yes

Administrators can use AWS 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, AWS evaluates them using a logical AND operation. If you specify multiple values for a single condition key, AWS 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*.

AWS supports global condition keys and service-specific condition keys. To see all AWS global condition keys, see [AWS global condition context keys]() in the *IAM User Guide*.

To see a list of Deadline Cloud condition keys, see [Condition keys for AWS Deadline Cloud]() in the *Service Authorization Reference*. To learn with which actions and resources you can use a condition key, see [Actions defined by AWS Deadline Cloud]().

To view examples of Deadline Cloud identity-based policies, see [Identity-based policy examples for Deadline Cloud]().

ACLs in Deadline Cloud

Supports ACLs: No

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

ABAC with Deadline Cloud

Supports ABAC (tags in policies): Yes

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. In AWS, these attributes are called *tags*. You can attach tags to IAM entities (users or roles) and to many AWS resources. Tagging entities and resources is the first step of ABAC. Then you design ABAC policies to allow operations when the principal's tag matches the tag on the resource that they are trying to access.

ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

To control access based on tags, you provide tag information in the [condition element](#) of a policy using the aws:ResourceTag/*key-name*, aws:RequestTag/*key-name*, or aws:TagKeys condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see [What is ABAC?](#) in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see [Use attribute-based access control (ABAC)](#) in the *IAM User Guide*.

Using temporary credentials with Deadline Cloud

Supports temporary credentials: Yes

Some AWS services don't work when you sign in using temporary credentials. For additional information, including which AWS services work with temporary credentials, see [AWS services that work with IAM](#) in the *IAM User Guide*.

You are using temporary credentials if you sign in to the AWS Management Console using any method except a user name and password. For example, when you access AWS using your company's single sign-on (SSO) link, that process automatically creates temporary credentials. You also automatically create temporary credentials when you sign in to the console as a user and then switch roles. For more information about switching roles, see [Switching to a role (console)](#) in the *IAM User Guide*.

You can manually create temporary credentials using the AWS CLI or AWS API. You can then use those temporary credentials to access AWS. AWS recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see [Temporary security credentials in IAM](#).

Forward access sessions for Deadline Cloud

Supports forward access sessions (FAS): Yes

When you use an IAM user or role to perform actions in AWS, 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 AWS service, combined with

the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS 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 roles for Deadline Cloud

Supports service roles: Yes

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 AWS service in the *IAM User Guide*.

> ⚠️ **Warning**
>
> Changing the permissions for a service role might break Deadline Cloud functionality. Edit service roles only when Deadline Cloud provides guidance to do so.

Service-linked roles for Deadline Cloud

Supports service-linked roles: No

A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For details about creating or managing service-linked roles, see AWS services that work with IAM. Find a service in the table that includes a Yes in the **Service-linked role** column. Choose the **Yes** link to view the service-linked role documentation for that service.

Identity-based policy examples for Deadline Cloud

By default, users and roles don't have permission to create or modify Deadline Cloud resources. They also can't perform tasks by using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS 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 Deadline Cloud, including the format of the ARNs for each of the resource types, see Actions, resources, and condition keys for AWS Deadline Cloud in the *Service Authorization Reference*.

Topics

- Policy best practices
- Using the Deadline Cloud console
- Policy to submit jobs to a queue
- Policy to allow creating a license endpoint
- Policy to allow monitoring a specific farm queue

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Deadline Cloud resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see AWS managed policies or AWS 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 AWS service, such as AWS 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 AWS 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 Deadline Cloud console

To access the AWS Deadline Cloud console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Deadline Cloud resources in your AWS 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.

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

To ensure that users and roles can still use the Deadline Cloud console, also attach the Deadline Cloud *ConsoleAccess* or *ReadOnly* AWS managed policy to the entities. For more information, see [Adding permissions to a user](#) in the *IAM User Guide*.

Policy to submit jobs to a queue

In this example, you create a scoped-down policy that grants permission to submit jobs to a specific queue in a specific farm.

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

```
            "Sid": "SubmitJobsFarmAndQueue",
            "Effect": "Allow",
            "Action": "deadline:CreateJob",
            "Resource": "arn:aws:deadline:REGION:ACCOUNT_ID:farm/FARM_A/queue/QUEUE_B/
 job/*"
        }
    ]
}
```

Policy to allow creating a license endpoint

In this example, you create a scoped-down policy that grants the required permissions to create and manage license endpoints. Use this policy to create the license endpoint for the VPC associated with your farm.

```
{
    "Version": "2012-10-17",
    "Statement": [{
        "SID": "CreateLicenseEndpoint",
        "Effect": "Allow",
        "Action": [
            "deadline:CreateLicenseEndpoint",
            "deadline:DeleteLicenseEndpoint",
            "deadline:GetLicenseEndpoint",
            "deadline:UpdateLicenseEndpoint",
            "deadline:ListLicenseEndpoints",
            "deadline:PutMeteredProduct",
            "deadline:DeleteMeteredProduct",
            "deadline:ListMeteredProducts",
            "deadline:ListAvailableMeteredProducts",
            "ec2:CreateVpcEndpoint",
            "ec2:DescribeVpcEndpoints",
            "ec2:DeleteVpcEndpoints"
        ],
        "Resource": "*"
    }]
}
```

Policy to allow monitoring a specific farm queue

In this example, you create a scoped-down policy that grants permission to monitor jobs in a specific queue for a specific farm.

```
{
    "Version": "2012-10-17",
    "Statement": [{
        "Sid": "MonitorJobsFarmAndQueue",
        "Effect": "Allow",
        "Action": [
            "deadline:SearchJobs",
            "deadline:ListJobs",
            "deadline:GetJob",
            "deadline:SearchSteps",
            "deadline:ListSteps",
            "deadline:ListStepConsumers",
            "deadline:ListStepDependencies",
            "deadline:GetStep",
            "deadline:SearchTasks",
            "deadline:ListTasks",
            "deadline:GetTask",
            "deadline:ListSessions",
            "deadline:GetSession",
            "deadline:ListSessionActions",
            "deadline:GetSessionAction"
        ],
        "Resource": [
            "arn:aws:deadline:REGION:123456789012:farm/FARM_A/queue/QUEUE_B",
            "arn:aws:deadline:REGION:123456789012:farm/FARM_A/queue/QUEUE_B/*"
        ]
    }]
}
```

AWS managed policies for Deadline Cloud

An AWS managed policy is a standalone policy that is created and administered by AWS. AWS 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 AWS managed policies might not grant least-privilege permissions for your specific use cases because they're available for all AWS 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 AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. AWS is most likely to update an AWS managed policy when a new AWS service is launched or new API operations become available for existing services.

For more information, see [AWS managed policies](#) in the *IAM User Guide*.

AWS managed policy: AWSDeadlineCloud-FleetWorker

You can attach the `AWSDeadlineCloud-FleetWorker` policy to your AWS Identity and Access Management (IAM) identities.

This policy grants workers in this fleet the permissions that are needed to connect to and receive tasks from the service.

Permissions details

This policy includes the following permissions:

- `deadline` – Allows principals to manage workers in a fleet.

For a JSON listing of the policy details, see [AWSDeadlineCloud-FleetWorker](#) in the *AWS Managed Policy reference guide*.

AWS managed policy: AWSDeadlineCloud-WorkerHost

You can attach the `AWSDeadlineCloud-WorkerHost` policy to your IAM identities.

This policy grants the permissions that are needed to initially connect to the service. It can be used as an Amazon Elastic Compute Cloud (Amazon EC2) instance profile.

Permissions details

This policy includes the following permissions:

- `deadline` – Allows principals to create workers.

For a JSON listing of the policy details, see [AWSDeadlineCloud-WorkerHost](#) in the *AWS Managed Policy reference guide*.

AWS managed policy: AWSDeadlineCloud-UserAccessFarms

You can attach the `AWSDeadlineCloud-UserAccessFarms` policy to your IAM identities.

This policy allows users to access farm data based on the farms that they are members of and their membership level.

Permissions details

This policy includes the following permissions:

- `deadline` – Allows the user to access farm data.
- `ec2` – Allows users to see details about Amazon EC2 instance types.
- `identitystore` – Allows users to see user and group names.

For a JSON listing of the policy details, see [AWSDeadlineCloud-UserAccessFarms](#) in the *AWS Managed Policy reference guide*.

AWS managed policy: AWSDeadlineCloud-UserAccessFleets

You can attach the `AWSDeadlineCloud-UserAccessFleets` policy to your IAM identities.

This policy allows users to access fleet data based on the farms that they are members of and their membership level.

Permissions details

This policy includes the following permissions:

- `deadline` – Allows the user to access farm data.
- `ec2` – Allows users to see details about Amazon EC2 instance types.
- `identitystore` – Allows users to see user and group names.

For a JSON listing of the policy details, see [AWSDeadlineCloud-UserAccessFleets](#) in the *AWS Managed Policy reference guide*.

AWS managed policy: AWSDeadlineCloud-UserAccessJobs

You can attach the `AWSDeadlineCloud-UserAccessJobs` policy to your IAM identities.

This policy allows users to access job data based on the farms that they are members of and their membership level.

Permissions details

This policy includes the following permissions:

- `deadline` – Allows the user to access farm data.
- `ec2` – Allows users to see details about Amazon EC2 instance types.
- `identitystore` – Allows users to see user and group names.

For a JSON listing of the policy details, see [AWSDeadlineCloud-UserAccessJobs](#) in the *AWS Managed Policy reference guide*.

AWS managed policy: AWSDeadlineCloud-UserAccessQueues

You can attach the `AWSDeadlineCloud-UserAccessQueues` policy to your IAM identities.

This policy allows users to access queue data based on the farms that they are members of and their membership level.

Permissions details

This policy includes the following permissions:

- `deadline` – Allows the user to access farm data.
- `ec2` – Allows users to see details about Amazon EC2 instance types.
- `identitystore` – Allows users to see user and group names.

For a JSON listing of the policy details, see [AWSDeadlineCloud-UserAccessQueues](#) in the *AWS Managed Policy reference guide*.

Deadline Cloud updates to AWS managed policies

View details about updates to AWS managed policies for Deadline Cloud since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the Deadline Cloud Document history page.

Change	Description	Date
AWSDeadlineCloud-UserAccessFarms – Change AWSDeadlineCloud-UserAccessJobs – Change AWSDeadlineCloud-UserAccessQueues – Change	Deadline Cloud added new actions `deadline:GetJobTemplate` and `deadline:ListJobParameterDefinitions` to allow you to resubmit jobs.	October 7, 2024
Deadline Cloud started tracking changes	Deadline Cloud started tracking changes to its AWS managed policies.	April 2, 2024

Troubleshooting AWS Deadline Cloud identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Deadline Cloud and IAM.

Topics

- I am not authorized to perform an action in Deadline Cloud
- I am not authorized to perform iam:PassRole
- I want to allow people outside of my AWS account to access my Deadline Cloud resources

I am not authorized to perform an action in Deadline Cloud

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 `deadline:`*GetWidget* permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
  deadline: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 `deadline:`*GetWidget* action.

If you need help, contact your AWS 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 Deadline Cloud.

Some AWS 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 Deadline Cloud. 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: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 AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my AWS account to access my Deadline Cloud 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 Deadline Cloud supports these features, see [How Deadline Cloud works with IAM](#).

- To learn how to provide access to your resources across AWS accounts that you own, see [Providing access to an IAM user in another AWS account that you own](#) in the *IAM User Guide*.

- To learn how to provide access to your resources to third-party AWS accounts, see [Providing access to AWS 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 [Cross account resource access in IAM](#) in the *IAM User Guide*.

Compliance validation for Deadline Cloud

To learn whether an AWS service is within the scope of specific compliance programs, see [AWS services in Scope by Compliance Program](#) and choose the compliance program that you are interested in. For general information, see [AWS Compliance Programs](#).

You can download third-party audit reports using AWS Artifact. For more information, see [Downloading Reports in AWS Artifact](#).

Your compliance responsibility when using AWS services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS 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 AWS that are security and compliance focused.

- [Architecting for HIPAA Security and Compliance on Amazon Web Services](#) – This whitepaper describes how companies can use AWS to create HIPAA-eligible applications.

> ⓘ **Note**
>
> Not all AWS services are HIPAA eligible. For more information, see the [HIPAA Eligible Services Reference](#).

- AWS Compliance Resources – This collection of workbooks and guides might apply to your industry and location.

- AWS Customer Compliance Guides – Understand the shared responsibility model through the lens of compliance. The guides summarize the best practices for securing AWS services and map the guidance to security controls across multiple frameworks (including National Institute of Standards and Technology (NIST), Payment Card Industry Security Standards Council (PCI), and International Organization for Standardization (ISO)).

- Evaluating Resources with Rules in the *AWS Config Developer Guide* – The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.

- AWS Security Hub – This AWS service provides a comprehensive view of your security state within AWS. Security Hub uses security controls to evaluate your AWS 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.

- Amazon GuardDuty – This AWS service detects potential threats to your AWS accounts, workloads, containers, and data by monitoring your environment for suspicious and malicious activities. GuardDuty can help you address various compliance requirements, like PCI DSS, by meeting intrusion detection requirements mandated by certain compliance frameworks.

- AWS Audit Manager – This AWS service helps you continuously audit your AWS usage to simplify how you manage risk and compliance with regulations and industry standards.

Resilience in Deadline Cloud

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure.

AWS Deadline Cloud does not back up data stored in your job attachments S3 bucket. You can enable backups of your job attachments data using any standard Amazon S3 backup mechanism, such as S3 Versioning or AWS Backup.

Infrastructure security in Deadline Cloud

As a managed service, AWS Deadline Cloud is protected by AWS global network security. For information about AWS security services and how AWS protects infrastructure, see AWS Cloud Security. To design your AWS environment using the best practices for infrastructure security, see Infrastructure Protection in *Security Pillar AWS Well-Architected Framework*.

You use AWS published API calls to access Deadline Cloud 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 AWS Security Token Service (AWS STS) to generate temporary security credentials to sign requests.

Deadline Cloud doesn't support using AWS PrivateLink virtual private cloud (VPC) endpoint policies. It uses the AWS PrivateLink default policy, which grants full access to the endpoint. For more information, see Default endpoint policy in the *AWS PrivateLink user guide*.

Configuration and vulnerability analysis in Deadline Cloud

AWS handles basic security tasks like guest operating system (OS) and database patching, firewall configuration, and disaster recovery. These procedures have been reviewed and certified by the appropriate third parties. For more details, see the following resources:

- Shared Responsibility Model

- Amazon Web Services: Overview of Security Processes (whitepaper)

AWS Deadline Cloud manages tasks on service-managed or customer-managed fleets:

- For service-managed fleets, Deadline Cloud manages the guest operating system.

- For customer-managed fleets, you are responsible for managing the operating system.

For additional information about configuration and vulnerability analysis for AWS Deadline Cloud, see

- [Security best practices for Deadline Cloud](#)

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 AWS, 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, AWS 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](#) and [aws:SourceAccount](#) global condition context keys in resource policies to limit the permissions that AWS Deadline Cloud 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 Amazon Resource Name (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:deadline:*: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 following example shows how you can use the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in Deadline Cloud to prevent the confused deputy problem.

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

```
    "Principal": {
      "Service": "deadline.amazonaws.com"
    },
    "Action": "deadline:ActionName",
    "Resource": [
      "*"
    ],
    "Condition": {
      "ArnLike": {
        "aws:SourceArn": "arn:aws:deadline:*:123456789012:*"
      },
      "StringEquals": {
        "aws:SourceAccount": "123456789012"
      }
    }
  }
}
```

Access AWS Deadline Cloud using an interface endpoint (AWS PrivateLink)

You can use AWS PrivateLink to create a private connection between your VPC and AWS Deadline Cloud. You can access Deadline Cloud as if it were in your VPC, without the use of an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Instances in your VPC don't need public IP addresses to access Deadline Cloud.

You establish this private connection by creating an *interface endpoint*, powered by AWS PrivateLink. We create an endpoint network interface in each subnet that you enable for the interface endpoint. These are requester-managed network interfaces that serve as the entry point for traffic destined for Deadline Cloud.

For more information, see Access AWS services through AWS PrivateLink in the *AWS PrivateLink Guide*.

Considerations for Deadline Cloud

Before you set up an interface endpoint for Deadline Cloud, see Access an AWS service using an interface VPC endpoint in the *AWS PrivateLink Guide*.

Deadline Cloud supports making calls to all of its API actions through the interface endpoint.

By default, full access to Deadline Cloud is allowed through the interface endpoint. Alternatively, you can associate a security group with the endpoint network interfaces to control traffic to Deadline Cloud through the interface endpoint.

Deadline Cloud doesn't support VPC endpoint policies. For more information, see Control access to VPC endpoints using endpoint policies in the *AWS PrivateLink Guide*.

Deadline Cloud endpoints

Deadline Cloud uses two endpoints for access to the service using AWS PrivateLink.

Workers use the com.amazonaws.*region*.deadline.scheduling endpoint to get tasks from the queue, report progress to Deadline Cloud, and to send task output back. If you are using a customer-managed fleet, the scheduling endpoint is the only endpoint that you need to create unless you are using management operations. For example, if a job creates more jobs, you need to enable the management endpoint to call the CreateJob operation.

The Deadline Cloud monitor uses the com.amazonaws.*region*.deadline.management to manage the resources in your farm, such as creating and modifying queues and fleets or getting lists of jobs, steps, and tasks.

Deadline Cloud also requires endpoints for the following AWS service endpoints:

- Deadline Cloud uses AWS STS to authenticate workers so that they can access job assets. For more information about AWS STS, see Temporary security credentials in IAM in the *AWS Identity and Access Management User Guide*.
- If you set up your customer-managed fleet in a subnet with no internet connection you must create a VPC endpoint for Amazon CloudWatch Logs so that workers can write logs. For more information, see Monitoring with CloudWatch.
- If you use job attachments, you must create a VPC endpoint for Amazon Simple Storage Service (Amazon S3) so that workers can access the attachments. For more information, see Job attachments in Deadline Cloud.

Create endpoints for Deadline Cloud

You can create interface endpoints for Deadline Cloud using either the Amazon VPC console or the AWS Command Line Interface (AWS CLI). For more information, see Create an interface endpoint in the *AWS PrivateLink Guide*.

Create management and scheduling endpoints for Deadline Cloud using the following service names. Replace *region* with the AWS Region where you've deployed Deadline Cloud.

```
com.amazonaws.region.deadline.management
```

```
com.amazonaws.region.deadline.scheduling
```

If you enable private DNS for the interface endpoints, you can make API requests to Deadline Cloud using its default Regional DNS name. For example, `worker.deadline.us-east-1.amazonaws.com` for worker operations, or `management.deadline.us-east-1.amazonaws.com` for all other operations.

You must also create an endpoint for AWS STS using the following service name:

```
com.amazonaws.region.sts
```

If your customer-managed fleet is on a subnet without an internet connection, you must create a CloudWatch Logs endpoint using the following service name:

```
com.amazonaws.region.logs
```

If you use job attachments to transfer files, you must create an Amazon S3 endpoint using the following service name:

```
com.amazonaws.region.s3
```

Security best practices for Deadline Cloud

AWS Deadline Cloud (Deadline Cloud) provides a number of security features to consider as you develop and implement your own security policies. The following best practices are general guidelines and don't represent a complete security solution. Because these best practices might not be appropriate or sufficient for your environment, treat them as helpful considerations rather than prescriptions.

> **ⓘ Note**
>
> For more information about the importance of many security topics, see the Shared
> Responsibility Model.

Data protection

For data protection purposes, we recommend that you protect AWS account credentials and set up
individual accounts with AWS 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 AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and
  securing personal data that is stored in Amazon Simple Storage Service (Amazon S3).
- If you require FIPS 140-2 validated cryptographic modules when accessing AWS 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 sensitive identifying information, such as your
customers' account numbers, into free-form fields such as a **Name** field. This includes when you
work with AWS Deadline Cloud or other AWS services using the console, API, AWS CLI, or AWS
SDKs. Any data that you enter into Deadline Cloud or other services might get picked up for
inclusion in diagnostic logs. When you provide a URL to an external server, don't include credentials
information in the URL to validate your request to that server.

AWS Identity and Access Management permissions

Manage access to AWS resources using users, AWS Identity and Access Management (IAM) roles,
and by granting the least privilege to users. Establish credential management policies and
procedures for creating, distributing, rotating, and revoking AWS access credentials. For more
information, see IAM Best Practices in the *IAM User Guide*.

Run jobs as users and groups

When using queue functionality in Deadline Cloud, it's a best practice to specify an operating system (OS) user and its primary group so that the OS user has least-privilege permissions for the queue's jobs.

When you specify a "Run as user" (and group), any processes for jobs submitted to the queue will be run using that OS user and will inherit that user's associated OS permissions.

The fleet and queue configurations combine to establish a security posture. On the queue side, the "Job run as user" and IAM role can be specified to use the OS and AWS permissions for the queue's jobs. The fleet defines the infrastructure (worker hosts, networks, mounted shared storage) that, when associated to a particular queue, run jobs within the queue. The data available on the worker hosts needs to be accessed by jobs from one or more associated queues. Specifying a user or group helps protect the data in jobs from other queues, other installed software, or other users with access to the worker hosts. When a queue is without a user, it runs as the agent user which can impersonate (sudo) any queue user. In this way, a queue without a user can escalate privileges to another queue.

Networking

To prevent traffic from being intercepted or redirected, it's essential to secure how and where your network traffic is routed.

We recommend that you secure your networking environment in the following ways:

- Secure Amazon Virtual Private Cloud (Amazon VPC) subnet route tables to control how IP layer traffic is routed.
- If you are using Amazon Route 53 (Route 53) as a DNS provider in your farm or workstation setup, secure access to the Route 53 API.
- If you connect to Deadline Cloud outside of AWS such as by using on-premises workstations or other data centers, secure any on-premises networking infrastructure. This includes DNS servers and route tables on routers, switches, and other networking devices.

Jobs and job data

Deadline Cloud jobs run within sessions on worker hosts. Each session runs one or more processes on the worker host, which generally require that you input data to produce output.

To secure this data, you can configure operating system users with queues. The worker agent uses the queue OS user to run session sub-processes. These sub-processes inherit the queue OS user's permissions.

We recommend that you follow best practices to secure access to the data these sub-processes access. For more information, see [Shared responsibility model](#).

Farm structure

You can arrange Deadline Cloud fleets and queues many ways. However, there are security implications with certain arrangements.

A farm has one of the most secure boundaries because it can't share Deadline Cloud resources with other farms, including fleets, queues, and storage profiles. However, you can share external AWS resources within a farm, which compromises the security boundary.

You can also establish security boundaries between queues within the same farm using the appropriate configuration.

Follow these best practices to create secure queues in the same farm:

- Associate a fleet only with queues within the same security boundary. Note the following:
  - After job runs on the worker host, data may remain behind, such as in a temporary directory or the queue user's home directory.
  - The same OS user runs all the jobs on a service-owned fleet worker host, regardless of which queue you submit the job to.
  - A job might leave processes running on a worker host, making it possible for jobs from other queues to observe other running processes.
- Ensure that only queues within the same security boundary share an Amazon S3 bucket for job attachments.
- Ensure that only queues within the same security boundary share an OS user.
- Secure any other AWS resources that are integrated into the farm to the boundary.

Job attachment queues

Job attachments are associated with a queue, which uses your Amazon S3 bucket.

- Job attachments write to and read from a root prefix in the Amazon S3 bucket. You specify this root prefix in the `CreateQueue` API call.

- The bucket has a corresponding `Queue Role`, which specifies the role that grants queue users access to the bucket and root prefix. When creating a queue, you specify the `Queue Role` Amazon Resource Name (ARN) alongside the job attachments bucket and root prefix.

- Authorized calls to the `AssumeQueueRoleForRead`, `AssumeQueueRoleForUser`, and `AssumeQueueRoleForWorker` API operations return a set of temporary security credentials for the `Queue Role`.

If you create a queue and reuse an Amazon S3 bucket and root prefix, there is a risk of information being disclosed to unauthorized parties. For example, QueueA and QueueB share the same bucket and root prefix. In a secure workflow, ArtistA has access to QueueA but not QueueB. However, when multiple queues share a bucket, ArtistA can access the data in QueueB data because it uses the same bucket and root prefix as QueueA.

The console sets up queues that are secure by default. Ensure that the queues have a distinct combination of Amazon S3 bucket and root prefix unless they're part of a common security boundary.

To isolate your queues, you must configure the `Queue Role` to only allow queue access to the bucket and root prefix. In the following example, replace each *placeholder* with your resource-specific information.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:GetObject",
        "s3:PutObject",
        "s3:ListBucket",
        "s3:GetBucketLocation"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:s3:::JOB_ATTACHMENTS_BUCKET_NAME",
        "arn:aws:s3:::JOB_ATTACHMENTS_BUCKET_NAME/JOB_ATTACHMENTS_ROOT_PREFIX/*"
      ],
      "Condition": {
```

```
            "StringEquals": { "aws:ResourceAccount": "ACCOUNT_ID" }
      }
    },
    {
      "Action": ["logs:GetLogEvents"],
      "Effect": "Allow",
      "Resource": "arn:aws:logs:REGION:ACCOUNT_ID:log-group:/aws/deadline/FARM_ID/*"
    }
  ]
}
```

You must also set a trust policy on the role. In the following example, replace the *placeholder* text with your resource-specific information.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": ["sts:AssumeRole"],
      "Effect": "Allow",
      "Principal": { "Service": "deadline.amazonaws.com" },
      "Condition": {
        "StringEquals": { "aws:SourceAccount": "ACCOUNT_ID" },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:deadline:REGION:ACCOUNT_ID:farm/FARM_ID"
        }
      }
    },
    {
      "Action": ["sts:AssumeRole"],
      "Effect": "Allow",
      "Principal": { "Service": "credentials.deadline.amazonaws.com" },
      "Condition": {
        "StringEquals": { "aws:SourceAccount": "ACCOUNT_ID" },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:deadline:REGION:ACCOUNT_ID:farm/FARM_ID"
        }
      }
    }
  ]
}
```

Custom software Amazon S3 buckets

You can add the following statement to your Queue Role to access custom software in your Amazon S3 bucket. In the following example, replace *SOFTWARE_BUCKET_NAME* with the name of your S3 bucket.

```
"Statement": [
    {
        "Action": [
            "s3:GetObject",
            "s3:ListBucket"
        ],
        "Effect": "Allow",
        "Resource": [
            "arn:aws:s3:::SOFTWARE_BUCKET_NAME",
            "arn:aws:s3:::SOFTWARE_BUCKET_NAME/*"
        ]
    }
]
```

For more information about Amazon S3 security best practices, see Security best practices for Amazon S3 in the *Amazon Simple Storage Service User Guide*.

Worker hosts

Secure worker hosts to help ensure that each user can only perform operations for their assigned role.

We recommend the following best practices to secure worker hosts:

- Don't use the same `jobRunAsUser` value with multiple queues unless jobs submitted to those queues are within the same security boundary.

- Don't set the queue `jobRunAsUser` to the name of the OS user that the worker agent runs as.

- Grant queue users least-privileged OS permissions required for the intended queue workloads. Ensure that they don't have filesystem write permissions to work agent program files or other shared software.

- Ensure only the root user on Linux and the `Administrator` owns account on Windows owns and can modify the worker agent program files.

- On Linux worker hosts, consider configuring a `umask` override in `/etc/sudoers` that allows the worker agent user to launch processes as queue users. This configuration helps ensure other users can't access files written to the queue.

- Grant trusted individuals least-privileged access to worker hosts.

- Restrict permissions to local DNS override configuration files (`/etc/hosts` on Linux and `C:\Windows\system32\etc\hosts` on Windows), and to route tables on workstations and worker host operating systems.

- Restrict permissions to DNS configuration on workstations and worker host operating systems.

- Regularly patch the operating system and all installed software. This approach includes software specifically used with Deadline Cloud such as submitters, adaptors, worker agents, OpenJD packages, and others.

- Use strong passwords for the Windows queue `jobRunAsUser`.

- Regularly rotate the passwords for your queue `jobRunAsUser`.

- Ensure least privilege access to the Windows password secretes and delete unused secrets.

- Don't give the queue `jobRunAsUser` permission the schedule commands to run in the future:

  - On Linux, deny these accounts access to `cron` and `at`.

  - On Windows, deny these accounts access to the Windows task scheduler.

> **ⓘ Note**
>
> For more information about the importance of regularly patching the operating system and installed software, see the Shared Responsibility Model.

Workstations

It's important to secure workstations with access to Deadline Cloud. This approach helps ensure that any jobs you submit to Deadline Cloud can't run arbitrary workloads billed to your AWS account.

We recommend the following best practice to secure artist workstations. For more information, see the Shared Responsibility Model.

- Secure any persisted credentials that provide access to AWS, including Deadline Cloud. For more information, see Managing access keys for IAM users in the *IAM User Guide*.

- Only install trusted, secure software.

- Require users federate with an identity provider to access AWS with temporary credentials.

- Use secure permissions on Deadline Cloud submitter program files to prevent tampering.

- Grant trusted individuals least-privileged access to artist workstations.

- Only use submitters and adaptors that you obtain through the Deadline Cloud Monitor.

- Restrict permissions to local DNS override configuration files (`/etc/hosts` on Linux and macOS, and `C:\Windows\system32\etc\hosts` on Windows), and to route tables on workstations and worker host operating systems.

- Restrict permissions to `/etc/resolve.conf` on workstations and worker host operating systems.

- Regularly patch the operating system and all installed software. This approach includes software specifically used with Deadline Cloud such as submitters, adaptors, worker agents, OpenJD packages, and others.

Document history

The following table describes important changes in each release of the *AWS Deadline Cloud Developer Guide*.

Change	Description	Date
[Create a conda package](#)	Added information about how to create a conda package for an application. For more information, see [Create a conda package](#).	August 29, 2024
[New guide](#)	This is the initial release of the *Deadline Cloud Developer Guide*.	July 26, 2024