

# AWS Batch

- [Tagging w/ Nextflow](#)
- [Troubleshooting for AWS Batch & Nextflow](#)

# Tagging w/ Nextflow

## Tagging ??

Tag (Tag) AWS resource. Tagging is a way to identify and categorize resources. Tags are key-value pairs that you can attach to an AWS resource. The tags and the values you provide are stored in your AWS account. Tags help you to manage your AWS resources more effectively. For example, you can use tags to track the cost of your resources, to organize your resources into groups, or to identify the owner of a resource.

AWS Batch is a managed service that allows you to run batch computing workloads on AWS. You can use the `aws batch` CLI to manage your batch jobs. The `aws batch` CLI has a `tags` command that you can use to tag your batch jobs. For example, you can use the following command to tag a batch job:

## AWS Batch

### Tagging your AWS Batch resources

Batch job definition (job definition) Propagate Tags AWS Batch job definition. The `Propagate Tags` parameter is a boolean value that indicates whether to propagate tags to the batch job. The default value is `false`. To propagate tags, you must set `Propagate Tags` to `true`.

AWS Batch jobs tag propagation. (Link : [https://docs.aws.amazon.com/batch/latest/userguide/using-tags.html?icmpid=docs\\_console\\_unmapped#tag-resources](https://docs.aws.amazon.com/batch/latest/userguide/using-tags.html?icmpid=docs_console_unmapped#tag-resources))

## Nextflow

Nextflow is a workflow tool that allows you to run Nextflow + AWS Batch. The `resourceLabels` parameter is a list of tags that you can use to tag your batch jobs. For example, you can use the following command to tag a batch job:

Nextflow is a workflow tool that allows you to run Nextflow + AWS Batch. The `resourceLabels` parameter is a list of tags that you can use to tag your batch jobs. (Link : <https://github.com/nextflow-io/nextflow/blob/5a37e6177f7a0e02b2af922768a0df5984b07b7b/plugins/nf-amazon/src/main/nextflow/cloud/aws/batch/AwsBatchTaskHandler.groovy#L713C36-L713C53>)

```

params.str = 'Hello world!'

process splitLetters {
  resourceLabels region: 'ap-northeast-2', user: 'brandon', project: 'nf-test'

  output:
  path 'chunk_*'

  """
  printf '${params.str}' | split -b 6 - chunk_
  """
}

process convertToUpper {
  resourceLabels region: 'ap-northeast-2', user: 'brandon', project: 'nf-test'

  input:
  path x

  output:
  stdout

  """
  cat $x | tr '[a-z]' '[A-Z]'
  """
}

workflow {
  splitLetters | flatten | convertToUpper | view { it.trim() }
}

```



 AWS Batch
 
 Job
 
 Tags
 


 .

Stopped at  
-

Key	Value
region	ap-northeast-2
project	nf-test
user	brandon

Key	Value
aws:batch:compute-environment	nextflow-ce
aws:batch:job-definition	nf-quay-io-nextflow-rnaseq-nf-v1-1:4
aws:batch:job-queue	nextflow-jq
myproject	innov_2024
nextflow.io/createdAt	2024-07-19T02:49:51.551Z
nextflow.io/version	24.04.3
project	nf-test
region	ap-northeast-2
user	brandon

# Troubleshooting for AWS Batch & Nextflow

## CannotPullImageManifestError: Error response from daemon: manifest unknown: manifest unknown

```
[ec2-user@ip-172-31-32-20 helloWorld]$ nextflow run tutorial.nf -bucket-dir s3:// /nextflow_tutorial_work/
```

```
NEXTFLOW ~ version 24.10.3
```

```
Launching `tutorial.nf` [mad_volhard] DSL2 - revision: e82728278a
```

```
executor > awsbatch (1)
[72/66dc9a] process > splitLetters [100%] 1 of 1, failed: 1 ✖
[-] process > convertToUpper -
ERROR ~ Error executing process > 'splitLetters'
```

```
Caused by:
Task failed to start - CannotPullImageManifestError: Error response from daemon: manifest unknown: manifest unknown
```

Command executed:

```
printf 'Hello world!' | split -b 6 - chunk_
```

Command exit status:

```
-
```

Command output:  
(empty)

Work dir:  
s3:// /nextflow\_tutorial\_work/72/66dc9afd5eced8b691b64fbf0e295d

Container:  
quay.io/biocontainers/salmon

Tip: when you have fixed the problem you can continue the execution adding the option `--resume` to the run command line

-- Check '.nextflow.log' file for details

Details

Status

🔴 Failed

Created at

Dec 26, 2024, 12:33:35 UTC+00:00

Status reason

Task failed to start

Started at

-

ARN

[arn:aws:batch:ap-northeast-2:985fb91-89ad-45e9-9aa1-75f73aea21e0](#):job/c985fb91-89ad-45e9-9aa1-75f73aea21e0

Stopped at

Dec 26, 2024, 12:37:56 UTC+00:00

Job ID

[c985fb91-89ad-45e9-9aa1-75f73aea21e0](#)

Log stream name

[nf-quay-io-biocontainers-salmon/default/7f01e52680b6482bb5efb855ff3d1a87](#)

Job attempts

Job configuration

Containers

Logging

Tags

JSON

Job attempts

Created at

Dec 26, 2024, 21:33:35

Started at

-

Stopped at

Dec 26, 2024, 21:37:56

Attempt #1

Status reason

Task failed to start

Started at

-

Stopped at

Dec 26, 2024, 21:37:56

Container Instance ARN

[arn:aws:ecs:ap-northeast-2:igib-a4723f19-7b71-35ba-8fb5-5ac4994eec97](#):container-instance/AWSBatch-IGIB-a4723f19-7b71-35ba-8fb5-5ac4994eec97/1bb2d81cb03f4382b8e0b5be4d0d5199

ECS task ARN

[arn:aws:ecs:ap-northeast-2:igib-a4723f19-7b71-35ba-8fb5-5ac4994eec97](#):task/AWSBatch-IGIB-a4723f19-7b71-35ba-8fb5-5ac4994eec97/7f01e52680b6482bb5efb855ff3d1a87

Exit code

-

Reason

CannotPullImageManifestError: Error response from daemon: manifest unknown: manifest unknown

Log stream name

[nf-quay-io-biocontainers-salmon/default/7f01e52680b6482bb5efb855ff3d1a87](#)

Network interfaces

Attachment ID	Ipv6 address	Private Ipv4 address
No resources		

## Solution

Check your docker image.

```
[ec2-user@ip-172-31-32-20 ~]$ docker pull quay.io/biocontainers/salmon
Using default tag: latest
Error response from daemon: manifest for quay.io/biocontainers/salmon:latest not found: manifest unknown: manifest unknown
```

bash: aws: command not found

```
[ec2-user@ip-172-31-32-20 helloWorld]$ nextflow run tutorial.nf -bucket-dir s3://.../nextflow_tutorial_work/

NEXTFLOW ~ version 24.10.3

Launching `tutorial.nf` [voluminous_pasteur] DSL2 - revision: e82728278a

executor > awsbatch (1)
[e6/c0cc8c] process > splitLetters [100%] 1 of 1, failed: 1 ✖
[-] process > convertToUpper -
ERROR ~ Error executing process > 'splitLetters'

Caused by:
  Essential container in task exited

Command executed:

  printf 'Hello world!' | split -b 6 - chunk_

Command exit status:
  127

Command output:
  (empty)

Command error:
  bash: aws: command not found
  bash: line 1: aws: command not found

Work dir:
  s3://.../nextflow_tutorial_work/e6/c0cc8c8f8eda05f0e821e094353adb

Container:
  quay.io/nextflow/rnaseq-nf:v1.1

Tip: when you have fixed the problem you can continue the execution adding the option `--resume` to the run command line
-- Check '.nextflow.log' file for details
```

## Solution

Check your container image has the aws cli program.

<https://www.nextflow.io/docs/latest/aws.html#custom-ami>

**ERROR ~ Pipeline failed. Please refer to troubleshooting docs: <https://nf-co.re/docs/usage/troubleshooting>**

```

Uploading local `bin` scripts folder to s3:// /nf-core_scrnaseq_workdir_2/tmp/b3/944d6c2946dc964109c3b6f42b5246/bin
-[nf-core/scrnaseq] Pipeline completed with errors-
executor > awsbatch (22)
[56/f664ed] process > NFCORE_SCRNASEQ:SCRNASEQ:FASTQC_CHECK:FASTQC (Sample_Y) [100%] 2 of 2 ✓
[23/3e00b7] process > NFCORE_SCRNASEQ:SCRNASEQ:CELLRANGER_ALIGN:CELLRANGER_COUNT (Sample_Y) [100%] 2 of 2 ✓
[3f/4b550e] process > NFCORE_SCRNASEQ:SCRNASEQ:MTX_TO_H5AD (Sample_Y) [100%] 4 of 4 ✓
[75/2c2dd9] process > NFCORE_SCRNASEQ:SCRNASEQ:H5AD_CONVERSION:CONCAT_H5AD (Sample_X) [75%] 3 of 4
[b1/f909b2] process > NFCORE_SCRNASEQ:SCRNASEQ:H5AD_CONVERSION:ANNDATAR_CONVERT (combined) [71%] 5 of 7
[f9/4df11a] process > NFCORE_SCRNASEQ:SCRNASEQ:EMPTY_DROPLET_REMOVAL:CELLBENDER_REMOVEBACKGROUND (Sample_Y) [0%] 0 of 2
[-] process > NFCORE_SCRNASEQ:SCRNASEQ:EMPTY_DROPLET_REMOVAL:ADATA_BARCODES -
[-] process > NFCORE_SCRNASEQ:SCRNASEQ:EMPTYDROPS_H5AD_CONVERSION:CONCAT_H5AD -
[-] process > NFCORE_SCRNASEQ:SCRNASEQ:EMPTYDROPS_H5AD_CONVERSION:ANNDATAR_CONVERT -
[6b/6e7032] process > NFCORE_SCRNASEQ:SCRNASEQ:MULTIQC [0%] 0 of 1
ERROR ~ Pipeline failed. Please refer to troubleshooting docs: https://nf-co.re/docs/usage/troubleshooting

-- Check '.nextflow.log' file for details
WARN: Killing running tasks (6)
WARN: Failed to render execution report -- see the log file for details
WARN: Failed to render execution timeline -- see the log file for details

```

## Solution

Check the .nextflow.log file content. There may be various causes, but this appears to be because the token has expired.

```

Caused by: com.amazonaws.services.s3.model.AmazonS3Exception: The provided token has expired. (Service: Amazon S3; Status Code: 400; Error Code: ExpiredToken; Request ID: SRHCCGA0PXY26G; S3 Extended Request ID: Sn4KngREFF0WVxzcVhSCYmXoxTNl0503rk6kgN05b8Fq+fnTNk/39rz2R6vASzoDLTyP31rj7I+; Proxy: n
ull)
at com.amazonaws.http.AmazonHttpClient$RequestExecutor.handleErrorResponse(AmazonHttpClient.java:1888)
at com.amazonaws.http.AmazonHttpClient$RequestExecutor.handleServiceErrorResponse(AmazonHttpClient.java:1418)
at com.amazonaws.http.AmazonHttpClient$RequestExecutor.executeOneRequest(AmazonHttpClient.java:1387)
at com.amazonaws.http.AmazonHttpClient$RequestExecutor.executeHelper(AmazonHttpClient.java:1157)
at com.amazonaws.http.AmazonHttpClient$RequestExecutor.doExecute(AmazonHttpClient.java:814)
at com.amazonaws.http.AmazonHttpClient$RequestExecutor.executeWithTimer(AmazonHttpClient.java:781)
at com.amazonaws.http.AmazonHttpClient$RequestExecutor.execute(AmazonHttpClient.java:755)
at com.amazonaws.http.AmazonHttpClient$RequestExecutor.access$500(AmazonHttpClient.java:715)
at com.amazonaws.http.AmazonHttpClient$RequestExecutor.runBuilderImpl(AmazonHttpClient.java:697)
at com.amazonaws.http.AmazonHttpClient.execute(AmazonHttpClient.java:561)
at com.amazonaws.http.AmazonHttpClient.execute(AmazonHttpClient.java:541)
at com.amazonaws.services.s3.AmazonS3Client.invoke(AmazonS3Client.java:5558)
at com.amazonaws.services.s3.AmazonS3Client.invoke(AmazonS3Client.java:5585)
at com.amazonaws.services.s3.AmazonS3Client.access$3$0(AmazonS3Client.java:423)
at com.amazonaws.services.s3.AmazonS3Client$PutObjectStrategy.invokeServiceCall(AmazonS3Client.java:6639)
at com.amazonaws.services.s3.AmazonS3Client.uploadObject(AmazonS3Client.java:1892)
at com.amazonaws.services.s3.AmazonS3Client.putObject(AmazonS3Client.java:1852)
at nextflow.cloud.aws.nio.S3OutputStream.putObject(S3OutputStream.java:636)
... 32 common frames omitted
Dec-28 13:52:05.072 [main] DEBUG nextflow.cache.CacheDB - Closing CacheDB done
Dec-28 13:52:05.081 [main] INFO org.pf4j.AbstractPluginManager - Stop plugin 'nf-amazon2.9.2'
Dec-28 13:52:05.081 [main] DEBUG nextflow.plugin.BasePlugin - Plugin stopped nf-amazon
Dec-28 13:52:05.081 [main] INFO org.pf4j.AbstractPluginManager - Stop plugin 'nf-schema2.1.1'
Dec-28 13:52:05.081 [main] DEBUG nextflow.plugin.BasePlugin - Plugin stopped nf-schema
Dec-28 13:52:05.083 [main] DEBUG nextflow.script.ScriptRunner - > Execution complete -- Goodbye
(END)

```