



Amazon EMR on EKS Development Guide

Amazon EMR



Amazon EMR: Amazon EMR on EKS Development Guide

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

What is Amazon EMR on EKS?	1
Architecture for Amazon EMR on EKS	2
Understanding Amazon EMR on EKS concepts and terminology	3
Kubernetes namespace	3
Virtual cluster	3
Job run	4
Amazon EMR containers	4
What happens when you submit work to an Amazon EMR on EKS virtual cluster	5
Getting started with Amazon EMR on EKS	6
Run a Spark application	7
Best practices	12
Security	12
Pyspark job submission	12
Storage	12
Metastore integration	13
Debugging	13
Troubleshooting Amazon EMR on EKS issues	13
Node placement	13
Performance	13
Cost optimization	14
Using AWS Outposts	14
Customizing Docker images	15
How to customize Docker images	15
Prerequisites	16
Step 1: Retrieve a base image from Amazon Elastic Container Registry (Amazon ECR)	16
Step 2: Customize a base image	17
Step 3: (Optional but recommended) Validate a custom image	18
Step 4: Publish a custom image	19
Step 5: Submit a Spark workload in Amazon EMR using a custom image	20
Customize Docker images for interactive endpoints	22
Work with multi-architecture images	24
Details for selecting a base image URI	26
Amazon ECR registry accounts	27
Considerations for customizing images	29

Running Flink jobs	31
Flink Kubernetes operator	31
Setting up	32
Installing the Flink Kubernetes operator	33
Run a Flink application	35
Security role permissions for running a Flink application	39
Uninstalling the operator	41
Flink Native Kubernetes	42
Setting up	42
Getting started	42
Security requirements	45
Customizing Docker images for Flink and FluentD	46
Prerequisites	46
Retrieve a base image from Amazon Elastic Container Registry	47
Customize a base image	47
Publish your custom image	48
Submit a Flink workload	49
Monitoring	49
Using Amazon Managed Service for Prometheus	50
Using the Flink UI	51
Using monitoring configuration	53
How Flink supports high availability and job resiliency	57
Using high availability	58
Optimizing restart times	64
Graceful decommission	70
Using Autoscaler	73
Autoscaler parameter autotuning	75
Maintenance and troubleshooting for Flink jobs on Amazon EMR on EKS	83
Maintaining Flink applications	84
Troubleshooting	85
Supported releases	88
Running Spark jobs	90
StartJobRun	90
Setting up	91
Submit a job run with StartJobRun	120
Using job submitter classification	123

Using Amazon EMR container defaults classification	129
Spark operator	134
Setting up	135
Getting started	136
Vertical autoscaling	140
Uninstall	145
Using monitoring configuration to monitor Spark	145
Security	152
spark-submit	163
Setting up	164
Getting started	165
Security	166
Apache Livy	171
Setting up	172
Getting started	173
Running a Spark application	178
Uninstalling	180
Security	180
Installation properties	190
Troubleshoot common environment-variable format errors	196
Managing job runs	197
Manage with CLI	197
Run Spark SQL scripts	203
Job run states	206
View jobs in the console	206
Common job run errors	207
Using job templates	213
Create and using a job template to start a job run	213
Defining job template parameters	215
Controlling access to job templates	217
Using pod templates	219
Common scenarios	219
Enabling pod templates with Amazon EMR on EKS	221
Pod template fields	223
Sidecar container considerations	226
Using retry policies	228

Set a retry policy	228
Retrieve the policy status	230
Monitor the job	231
Find the driver logs	232
Using Spark event log rotation	232
Using Spark container log rotation	233
Using vertical autoscaling	235
Setting up	236
Getting started	238
Configuration	240
Monitoring the recommendations	246
Uninstalling	247
Running interactive workloads	248
Overview of interactive endpoints	248
Interactive endpoints prerequisites	250
AWS CLI	250
eksctl	250
Amazon EKS cluster	251
Grant Cluster access	251
Activate IAM roles for Service Accounts	251
Create IAM job execution role	252
Grant users access	252
Register Amazon EKS cluster with Amazon EMR	252
Load Balancer Controller	253
Creating an interactive endpoint	253
Create an interactive endpoint	253
Specify custom parameters	254
.....	255
Interactive endpoint parameters	255
Configuring settings for interactive endpoints	257
Monitoring Spark jobs	257
Custom pod templates	258
Deploying a JEG pod to a node group	259
JEG configuration options	263
Modifying PySpark parameters	263
Custom kernel image	264

Monitoring interactive endpoints	266
Examples	268
Using self-hosted Jupyter notebooks	269
Create a security group	269
Create an interactive endpoint	270
Get the gateway server URL	270
Get the auth token	271
Deploy the notebook	272
Clean up	277
Getting information about interactive endpoints with CLI commands	278
.....	278
List interactive endpoints	279
Delete interactive endpoint	281
Uploading data	282
Prerequisites	282
Getting started	282
Monitoring jobs	285
Monitor jobs with Amazon CloudWatch Events	285
Automate Amazon EMR on EKS with CloudWatch Events	286
Example: Set up a rule that invokes Lambda	287
Monitor job's driver pod with a retry policy using Amazon CloudWatch Events	288
Managing virtual clusters	289
Create a virtual cluster	289
List virtual clusters	291
Describe a virtual cluster	291
Delete a virtual cluster	291
Virtual cluster states	291
Concurrent job limits for virtual clusters	292
Key benefits of concurrent job limits	292
Getting started with concurrent job limits	293
Choosing values for maxConcurrentJobRuns and maxInQueueJobRuns	294
Considerations for concurrent job limits	295
Tutorials	296
Using Delta Lake	296
Using Iceberg	297
Spark session configurations for catalog integration	298

Using PyFlink	299
Using AWS Glue with Flink	300
Using Apache Hudi	303
Submit an Apache Hudi job	303
Using Spark cuDF	307
Using Spark on Redshift	311
Launch a Spark application	311
Authenticate to Amazon Redshift	312
Read and write to Amazon Redshift	315
Considerations	317
Using Volcano	318
Overview	318
Installation	318
Submit: Spark operator	319
Submit: spark-submit	321
Using YuniKorn	322
Overview	323
Create your cluster	323
Install YuniKorn	325
Submit: Spark operator	325
Submit: spark-submit	328
Security	12
Best practices	332
Apply principle of least privilege	332
Access control list for endpoints	332
Get the latest security updates for custom images	332
Limit pod credential access	333
Isolate untrusted application code	333
Role-based access control (RBAC) permissions	333
Restrict access to nodegroup IAM role or instance profile credentials	334
Data protection	334
Encryption at rest	335
Encryption in transit	338
Identity and Access Management	338
Audience	339
Authenticating with identities	339

Managing access using policies	340
How Amazon EMR on EKS works with IAM	342
Using Service-Linked Roles	347
Managed policies for Amazon EMR on EKS	351
Using job execution roles with Amazon EMR on EKS	352
Identity-based policy examples	355
Policies for tag-based access control	358
Troubleshooting	361
Using Amazon EMR on EKS with AWS Lake Formation	363
How Amazon EMR on EKS works with AWS Lake Formation	364
Enable Lake Formation with Amazon EMR on EKS	365
Considerations and limitations	373
Troubleshooting	376
Logging and monitoring	378
Encrypting logs	378
CloudTrail logs	382
S3 Access Grants	384
Overview	384
Launch a cluster	385
Considerations	386
Compliance Validation	386
Resilience	387
Infrastructure Security	387
Configuration and vulnerability analysis	387
Interface VPC endpoints	388
Create a VPC Endpoint Policy for Amazon EMR on EKS	389
Cross-account access	392
Prerequisites	392
How to access a cross-account Amazon S3 bucket or DynamoDB table	392
Tagging resources	397
Tag basics	397
Tag your resources	398
Tag restrictions	399
Work with tags using the AWS CLI and the Amazon EMR on EKS API	399
Troubleshooting	13
PVC job failures	401

Verification	401
Patch	402
Manual patch	405
Vertical autoscaling failures	408
403 Forbidden error	408
Namespace not found	408
Docker credentials error	408
Spark operator failures	409
Helm chart install failed	409
Unsupported filesystem exception	409
Service endpoints and quotas	411
Service endpoints	411
Service quotas	413
View quotas and request quota increases	415
Release versions	416
emr-spark-8.0.0 releases	417
Releases	418
Release notes	418
Changes and features	420
emr-spark-8.0.0-latest	420
emr-spark-8.0.0-20260421	420
7.13.0 releases	421
Releases	421
Release notes	423
Changes and features	424
emr-7.13.0-latest	424
emr-7.13.0-20260410	424
emr-7.13.0 common vulnerabilities and exposures	425
emr-7.13.0-flink-latest	425
emr-7.13.0-flink-20260410	425
7.12.0 releases	426
Releases	426
Release notes	427
Changes and features	429
emr-7.12.0-latest	429
emr-7.12.0-20251111	430

emr-7.12.0 common vulnerabilities and exposures	430
emr-7.12.0-flink-latest	431
emr-7.12.0-flink-20251111	431
7.11.0 releases	431
Releases	431
Release notes	433
Changes and features	435
emr-7.11.0-latest	435
emr-7.11.0-20251020	435
emr-7.11.0-flink-latest	435
emr-7.11.0-flink-20251020	436
7.10.0 releases	436
Releases	436
Release notes	438
Changes and features	439
emr-7.10.0-latest	439
emr-7.10.0-20250801	440
emr-7.10.0 common vulnerabilities and exposures	440
emr-7.10.0-flink-latest	444
emr-7.10.0-flink-20250801	444
7.9.0 releases	445
Releases	445
Release notes	446
Changes	448
emr-7.9.0-latest	448
emr-7.9.0-20250425	448
emr-7.9.0 common vulnerabilities and exposures	449
emr-7.9.0-flink-latest	453
emr-7.9.0-flink-20250425	453
7.8.0 releases	453
Releases	453
Release notes	455
Changes	456
emr-7.8.0-latest	457
emr-7.8.0-20250228	457
emr-7.8.0-flink-latest	457

emr-7.8.0-flink-20250228	457
7.7.0 releases	458
Releases	458
Release notes	459
Changes	461
emr-7.7.0-latest	461
emr-7.7.0-20250131	461
emr-7.7.0-flink-latest	462
emr-7.7.0-flink-20250131	462
7.6.0 releases	462
Releases	462
Release notes	464
Features	465
Changes	466
emr-7.6.0-latest	466
emr-7.6.0-20241213	466
emr-7.6.0-flink-latest	466
emr-7.6.0-flink-20241213	467
7.5.0 releases	467
Releases	467
Release notes	467
emr-7.5.0 common vulnerabilities and exposures	467
7.4.0 releases	468
Releases	468
Release notes	468
7.3.0 releases	469
Releases	469
Release notes	470
Features	472
Changes	472
emr-7.3.0-latest	473
emr-7.3.0-20240920	473
emr-7.3.0 common vulnerabilities and exposures	473
emr-7.3.0-flink-latest	476
emr-7.3.0-flink-29240920	476
7.2.0 releases	476

Releases	476
Release notes	478
Features	479
emr-7.2.0-latest	480
emr-7.2.0-20240610	480
emr-7.2.0 common vulnerabilities and exposures	481
emr-7.2.0-flink-latest	484
emr-7.2.0-flink-20240610	484
7.1.0 releases	485
Releases	485
Release notes	486
Features	488
emr-7.1.0-latest	488
emr-7.1.0-20240321	489
emr-7.1.0 common vulnerabilities and exposures	489
emr-7.1.0-flink-latest	489
emr-7.1.0-flink-20240321	489
7.0.0 releases	490
Releases	490
Release notes	491
Features	493
Changes	493
emr-7.0.0-latest	493
emr-7.0.0-2024321	494
emr-7.0.0-20231211	494
emr-7.0.0 common vulnerabilities and exposures	494
emr-7.0.0-flink-latest	496
emr-7.0.0-flink-2024321	496
emr-7.0.0-flink-20231211	497
6.15.0 releases	497
Releases	497
Release notes	499
Features	500
emr-6.15.0-latest	500
emr-6.15.0-20240105	501
emr-6.15.0-20231109	501

emr-6.15.0-flink-latest	501
emr-6.15.0-flink-20240105	501
emr-6.15.0-flink-20231109	502
6.14.0 releases	502
Releases	502
Release notes	503
Features	505
emr-6.14.0-latest	505
emr-6.14.0-20231005	505
6.13.0 releases	505
Releases	506
Release notes	507
Features	508
emr-6.13.0-latest	509
emr-6.13.0-20230814	509
6.12.0 releases	509
Releases	509
Release notes	510
Features	512
emr-6.12.0-latest	512
emr-6.12.0-20240321	512
emr-6.12.0-20230701	512
6.11.0 releases	513
Releases	513
Release notes	513
Features	515
emr-6.11.0-latest	515
emr-6.11.0-20230905	516
emr-6.11.0-20230509	516
6.10.0 releases	516
emr-6.10.0-latest	519
emr-6.10.0-20230905	519
emr-6.10.0-20230624	519
emr-6.10.0-20230421	519
emr-6.10.0-20230403	520
emr-6.10.0-20230220	520

6.9.0 releases	520
emr-6.9.0-latest	523
emr-6.9.0-20230905	523
emr-6.9.0-20230624	523
emr-6.9.0-20221108	524
6.8.0 releases	524
emr-6.8.0-latest	528
emr-6.8.0-20230905	528
emr-6.8.0-20230624	528
emr-6.8.0-20221219	528
emr-6.8.0-20220802	529
6.7.0 releases	529
emr-6.7.0-latest	531
emr-6.7.0-20240321	531
emr-6.7.0-20230624	531
emr-6.7.0-20221219	531
emr-6.7.0-20220630	532
6.6.0 releases	532
emr-6.6.0-latest	533
emr-6.6.0-20240321	534
emr-6.6.0-20230624	534
emr-6.6.0-20221219	534
emr-6.6.0-20220411	534
6.5.0 releases	535
emr-6.5.0-latest	536
emr-6.5.0-20240321	536
emr-6.5.0-20221219	536
emr-6.5.0-20220802	537
emr-6.5.0-20211119	537
6.4.0 releases	537
emr-6.4.0-latest	538
emr-6.4.0-20240321	539
emr-6.4.0-20221219	539
emr-6.4.0-20210830	539
6.3.0 releases	539
emr-6.3.0-latest	541

emr-6.3.0-20240321	541
emr-6.3.0-20220802	541
emr-6.3.0-20211008	541
emr-6.3.0-20210802	542
emr-6.3.0-20210429	542
6.2.0 releases	542
emr-6.2.0-latest	543
emr-6.2.0-20240321	544
emr-6.2.0-20220802	544
emr-6.2.0-20211008	544
emr-6.2.0-20210802	544
emr-6.2.0-20210615	545
emr-6.2.0-20210129	545
emr-6.2.0-20201218	545
emr-6.2.0-20201201	545
5.36.0 releases	546
emr-5.36.0-latest	547
emr-5.36.0-20240321	547
emr-5.36.0-20221219	547
emr-5.36.0-20220620	547
emr-5.36.0-20220525	548
5.35.0 releases	548
emr-5.35.0-latest	549
emr-5.35.0-20240321	549
emr-5.35.0-20221219	550
emr-5.35.0-20220802	550
emr-5.35.0-20220307	550
5.34 releases	550
emr-5.34.0-latest	551
emr-5.34.0-20240321	552
emr-5.34.0-20220802	552
emr-5.34.0-20211208	552
5.33.0 releases	552
emr-5.33.0-latest	554
emr-5.33.0-20240321	554
emr-5.33.0-20221219	554

emr-5.33.0-20220802	554
emr-5.33.0-20211008	555
emr-5.33.0-20210802	555
emr-5.33.0-20210615	555
emr-5.33.0-20210323	555
5.32.0 releases	556
emr-5.32.0-latest	557
emr-5.32.0-20240321	557
emr-5.32.0-20220802	557
emr-5.32.0-20211008	558
emr-5.32.0-20210802	558
emr-5.32.0-20210615	558
emr-5.32.0-20210129	558
emr-5.32.0-20201218	559
emr-5.32.0-20201201	559
Document history	560

What is Amazon EMR on EKS?


Amazon EMR on EKS provides a deployment option for Amazon EMR that allows you to run open-source big data frameworks on Amazon Elastic Kubernetes Service (Amazon EKS). With this deployment option, you can focus on running analytics workloads while Amazon EMR on EKS builds, configures, and manages containers for open-source applications.

If you already use Amazon EMR, you can now run Amazon EMR based applications with other types of applications on the same Amazon EKS cluster. This deployment option also improves resource utilization and simplifies infrastructure management across multiple Availability Zones. If you already run big data frameworks on Amazon EKS, you can now use Amazon EMR to automate provisioning and management, and run Apache Spark more quickly.

Amazon EMR on EKS enables your team to collaborate more efficiently and process vast amounts of data more easily and cost-effectively:

- You can run applications on a common pool of resources without having to provision infrastructure. You can use [Amazon EMR Studio](#) and the AWS SDK or AWS CLI to develop, submit, and diagnose analytics applications running on EKS clusters. You can run scheduled jobs on Amazon EMR on EKS using self-managed Apache Airflow or Amazon Managed Workflows for Apache Airflow (MWAA).
- Infrastructure teams can centrally manage a common computing platform to consolidate Amazon EMR workloads with other container-based applications. You can simplify infrastructure management with common Amazon EKS tools and take advantage of a shared cluster for workloads that need different versions of open-source frameworks. You can also reduce operational overhead with automated Kubernetes cluster management and OS patching. With Amazon EC2 and AWS Fargate, you can enable multiple compute resources to meet performance, operational, or financial requirements.

The following diagram shows the two different deployment models for Amazon EMR.



Topics


- [Architecture for Amazon EMR on EKS](#)
- [Understanding Amazon EMR on EKS concepts and terminology](#)
- [What happens when you submit work to an Amazon EMR on EKS virtual cluster](#)

Architecture for Amazon EMR on EKS

Amazon EMR on EKS loosely couples applications to the infrastructure that they run on. Each infrastructure layer provides orchestration for the subsequent layer. When you submit a job to Amazon EMR, your job definition contains all of its application-specific parameters. Amazon EMR uses these parameters to instruct Amazon EKS about which pods and containers to deploy. Amazon EKS then brings online the computing resources from Amazon EC2 and AWS Fargate required to run the job.

With this loose coupling of services, you can run multiple, securely isolated jobs simultaneously. You can also benchmark the same job with different compute backends or spread your job across multiple Availability Zones to improve availability.

The following diagram illustrates how Amazon EMR on EKS works with other AWS services.



Understanding Amazon EMR on EKS concepts and terminology

Amazon EMR on EKS provides a deployment option for Amazon EMR that allows you to run open-source big data frameworks on Amazon Elastic Kubernetes Service (Amazon EKS). This topic gives you context on some of the common terminology for it, including namespaces, virtual clusters, and job runs, which are units of work that you submit for processing.

Kubernetes namespace

Amazon EKS uses Kubernetes namespaces to divide cluster resources between multiple users and applications. These namespaces are the foundation for multi-tenant environments. A Kubernetes namespace can have either Amazon EC2 or AWS Fargate as the compute provider. This flexibility provides you with different performance and cost options for your jobs to run on.

Virtual cluster

A virtual cluster is a Kubernetes namespace that Amazon EMR is registered with. Amazon EMR uses virtual clusters to run jobs and host endpoints. Multiple virtual clusters can be backed by the same

physical cluster. However, each virtual cluster maps to one namespace on an EKS cluster. Virtual clusters do not create any active resources that contribute to your bill or that require lifecycle management outside the service.

Job run

A job run is a unit of work, such as a Spark jar, PySpark script, or SparkSQL query, that you submit to Amazon EMR on EKS. One job can have multiple job runs. When you submit a job run, you include the following information:

- A virtual cluster where the job should run.
- A job name to identify the job.
- The execution role — a scoped IAM role that runs the job and allows you to specify which resources can be accessed by the job.
- The Amazon EMR release label that specifies the version of open-source applications to use.
- The artifacts to use when submitting your job, such as spark-submit parameters.

By default, logs are uploaded to the Spark History server and are accessible from the AWS Management Console. You can also push event logs, execution logs, and metrics to Amazon S3 and Amazon CloudWatch.

Amazon EMR containers

Amazon EMR containers is the [API name for Amazon EMR on EKS](#). The `emr-containers` prefix is used in the following scenarios:


- It is the prefix in the CLI commands for Amazon EMR on EKS. For example, `aws emr-containers start-job-run`.
- It is the prefix before IAM policy actions for Amazon EMR on EKS. For example, "Action": ["emr-containers:StartJobRun"]. For more information, see [Policy actions for Amazon EMR on EKS](#).
- It is the prefix used in Amazon EMR on EKS service endpoints. For example, `emr-containers.us-east-1.amazonaws.com`. For more information, see [Amazon EMR on EKS Service Endpoints](#).

What happens when you submit work to an Amazon EMR on EKS virtual cluster

Registering Amazon EMR with a Kubernetes namespace on Amazon EKS creates a virtual cluster. Amazon EMR can then run analytics workloads on that namespace. When you use Amazon EMR on EKS to submit Spark jobs to the virtual cluster, Amazon EMR on EKS requests the Kubernetes scheduler on Amazon EKS to schedule pods.

The following steps and diagram illustrate the Amazon EMR on EKS workflow:

- Use an existing Amazon EKS cluster or create one by using the [eksctl](#) command line utility or Amazon EKS console.
- Create a virtual cluster by registering Amazon EMR with a namespace on an EKS cluster.
- Submit your job to the virtual cluster using the AWS CLI or SDK.



For each job that you run, Amazon EMR on EKS creates a container with an Amazon Linux 2 base image, Apache Spark, and associated dependencies. Each job runs in a pod that downloads the container and starts to run it. The pod terminates after the job terminates. If the container's image has been previously deployed to the node, then a cached image is used and the download is bypassed. Sidecar containers, such as log or metric forwarders, can be deployed to the pod. After the job terminates, you can still debug it using Spark application UI in the Amazon EMR console.

Getting started with Amazon EMR on EKS

This topic helps you get started using Amazon EMR on EKS by deploying a Spark application on a virtual cluster. It includes steps to set up the correct permissions and to start a job. Before you begin, make sure that you've completed the steps in [Setting up Amazon EMR on EKS](#). This helps you get tools like the AWS CLI set up prior to creating your virtual cluster. For other templates that can help you get started, see our [EMR Containers Best Practices Guide](#) on GitHub.

You will need the following information from the setup steps:

- Virtual cluster ID for the Amazon EKS cluster and Kubernetes namespace registered with Amazon EMR

Important

When creating an EKS cluster, make sure to use m5.xlarge as the instance type, or any other instance type with a higher CPU and memory. Using an instance type with lower CPU or memory than m5.xlarge may lead to job failure due to insufficient resources available in the cluster.

- Name of the IAM role used for job execution
- Release label for the Amazon EMR release (for example, emr-6.4.0-latest)
- Destination targets for logging and monitoring:
 - Amazon CloudWatch log group name and log stream prefix
 - Amazon S3 location to store event and container logs

Important

Amazon EMR on EKS jobs use Amazon CloudWatch and Amazon S3 as destination targets for monitoring and logging. You can monitor job progress and troubleshoot failures by viewing the job logs sent to these destinations. To enable logging, the IAM policy associated with the IAM role for job execution must have the required permissions to access the target resources. If the IAM policy doesn't have the required permissions, you must follow the steps outlined in [Update the trust policy of the job execution role](#), [Configure](#)

[a job run to use Amazon S3 logs](#), and [Configure a job run to use CloudWatch Logs](#) before running this sample job.

Run a Spark application

Take the following steps to run a simple Spark application on Amazon EMR on EKS. The application entryPoint file for a Spark Python application is located at `s3://REGION.elasticmapreduce/emr-containers/samples/wordcount/scripts/wordcount.py`. The *REGION* is the Region in which your Amazon EMR on EKS virtual cluster resides, such as *us-east-1*.

1. Update the IAM policy for the job execution role with the required permissions, as the following policy statements demonstrate.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ReadFromLoggingAndInputScriptBuckets",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3::*.elasticmapreduce",
        "arn:aws:s3::*.elasticmapreduce/*",
        "arn:aws:s3:::amzn-s3-demo-bucket",
        "arn:aws:s3:::amzn-s3-demo-bucket/*",
        "arn:aws:s3:::amzn-s3-demo-bucket-b",
        "arn:aws:s3:::amzn-s3-demo-bucket-b/*"
      ]
    },
    {
      "Sid": "WriteToLoggingAndOutputDataBuckets",
      "Effect": "Allow",
      "Action": [
```



```

        "s3:PutObject",
        "s3:DeleteObject"
    ],
    "Resource": [
        "arn:aws:s3:::amzn-s3-demo-bucket/*",
        "arn:aws:s3:::amzn-s3-demo-bucket-b/*"
    ]
},
{
    "Sid": "DescribeAndCreateCloudwatchLogStream",
    "Effect": "Allow",
    "Action": [
        "logs:CreateLogStream",
        "logs:DescribeLogGroups",
        "logs:DescribeLogStreams"
    ],
    "Resource": [
        "arn:aws:logs:*:*:*"
    ]
},
{
    "Sid": "WriteToCloudwatchLogs",
    "Effect": "Allow",
    "Action": [
        "logs:PutLogEvents"
    ],
    "Resource": [
        "arn:aws:logs:*:*:log-group:my_log_group_name:log-
stream:my_log_stream_prefix/*"
    ]
}
]
}

```

- The first statement `ReadFromLoggingAndInputScriptBuckets` in this policy grants `ListBucket` and `GetObjects` access to the following Amazon S3 buckets:
 - *REGION.elasticmapreduce* - the bucket where the application `entryPoint` file is located.
 - *amzn-s3-demo-destination-bucket* - a bucket that you define for your output data.
 - *amzn-s3-demo-logging-bucket* - a bucket that you define for your logging data.

- The second statement `WriteToLoggingAndOutputDataBuckets` in this policy grants the job permissions to write data to your output and logging buckets respectively.
 - The third statement `DescribeAndCreateCloudwatchLogStream` grants the job with permissions to describe and create Amazon CloudWatch Logs.
 - The fourth statement `WriteToCloudwatchLogs` grants permissions to write logs to an Amazon CloudWatch log group named *my_log_group_name* under a log stream named *my_log_stream_prefix*.
2. To run a Spark Python application, use the following command. Replace all the replaceable *red italicized* values with appropriate values. The *REGION* is the Region in which your Amazon EMR on EKS virtual cluster resides, such as *us-east-1*.

```
aws emr-containers start-job-run \
--virtual-cluster-id cluster_id \
--name sample-job-name \
--execution-role-arn execution-role-arn \
--release-label emr-6.4.0-latest \
--job-driver '{
  "sparkSubmitJobDriver": {
    "entryPoint": "s3://REGION.elasticmapreduce/emr-containers/samples/wordcount/
scripts/wordcount.py",
    "entryPointArguments": ["s3://amzn-s3-demo-destination-bucket/
wordcount_output"],
    "sparkSubmitParameters": "--conf spark.executor.instances=2 --
conf spark.executor.memory=2G --conf spark.executor.cores=2 --conf
spark.driver.cores=1"
  }
}' \
--configuration-overrides '{
  "monitoringConfiguration": {
    "cloudWatchMonitoringConfiguration": {
      "logGroupName": "my_log_group_name",
      "logStreamNamePrefix": "my_log_stream_prefix"
    },
    "s3MonitoringConfiguration": {
      "logUri": "s3://amzn-s3-demo-logging-bucket"
    }
  }
}'
```

The output data from this job will be available at `s3://amzn-s3-demo-destination-bucket/wordcount_output`.

You can also create a JSON file with specified parameters for your job run. Then run the `start-job-run` command with a path to the JSON file. For more information, see [Submit a job run with StartJobRun](#). For more details about configuring job run parameters, see [Options for configuring a job run](#).

3. To run a Spark SQL application, use the following command. Replace all the *red italicized* values with appropriate values. The *REGION* is the Region in which your Amazon EMR on EKS virtual cluster resides, such as *us-east-1*.

```
aws emr-containers start-job-run \
--virtual-cluster-id cluster_id \
--name sample-job-name \
--execution-role-arn execution-role-arn \
--release-label emr-6.7.0-latest \
--job-driver '{
  "sparkSqlJobDriver": {
    "entryPoint": "s3://query-file.sql",
    "sparkSqlParameters": "--conf spark.executor.instances=2 --
conf spark.executor.memory=2G --conf spark.executor.cores=2 --conf
spark.driver.cores=1"
  }
}' \
--configuration-overrides '{
  "monitoringConfiguration": {
    "cloudWatchMonitoringConfiguration": {
      "logGroupName": "my_log_group_name",
      "logStreamNamePrefix": "my_log_stream_prefix"
    },
    "s3MonitoringConfiguration": {
      "logUri": "s3://amzn-s3-demo-logging-bucket"
    }
  }
}'
```

A sample SQL query file is shown below. You must have an external file store, such as S3, where the data for the tables is stored.

```
CREATE DATABASE demo;
```


```
CREATE EXTERNAL TABLE IF NOT EXISTS demo.amazonreview( marketplace string,
customer_id string, review_id string, product_id string, product_parent string,
product_title string, star_rating integer, helpful_votes integer, total_votes
integer, vine string, verified_purchase string, review_headline string,
review_body string, review_date date, year integer) STORED AS PARQUET LOCATION
's3://URI to parquet files';
SELECT count(*) FROM demo.amazonreview;
SELECT count(*) FROM demo.amazonreview WHERE star_rating = 3;
```

The output for this job will be available in the driver's stdout logs in S3 or CloudWatch, depending on the `monitoringConfiguration` that is configured.

4. You can also create a JSON file with specified parameters for your job run. Then run the `start-job-run` command with a path to the JSON file. For more information, see [Submit a job run](#). For more details about configuring job run parameters, see [Options for configuring a job run](#).

To monitor the progress of the job or to debug failures, you can inspect logs uploaded to Amazon S3, CloudWatch Logs, or both. Refer to log path in Amazon S3 at [Configure a job run to use S3 logs](#) and for Cloudwatch logs at [Configure a job run to use CloudWatch Logs](#). To see logs in CloudWatch Logs, follow the instructions below.

- Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
- In the **Navigation** pane, choose **Logs**. Then choose **Log groups**.
- Choose the log group for Amazon EMR on EKS and then view the uploaded log events.



Important

Jobs have a [default configured retry policy](#). For information on how to modify or disable the configuration, refer to [Using job retry policies](#).

Links to Amazon EMR on EKS best practices guides on GitHub

We've built the [Amazon EMR on EKS Best Practices Guide](#) using open source community collaboration so that we can iterate quickly and provide recommendations for aspects of creating and running a virtual cluster. We recommend that you use the [Amazon EMR on EKS best practices guide](#) for the sections. Choose the links in each section to go to the GitHub site.

Security

Note

For more information on security with Amazon EMR on EKS, see [Amazon EMR on EKS security best practices](#).

[Encryption best practices](#): how to use encryption for data at rest and in transit.

[Managing network security](#) describes how to configure security groups for pods for Amazon EMR on EKS while you connect to data sources that are hosted in AWS services like Amazon RDS and Amazon Redshift.

[Using AWS secrets manager to store secrets](#).

Pyspark job submission

[Pyspark job submission](#): specifies different types of packaging for pySpark applications using packaging formats like zip, egg, wheel, and pex.

Storage

[Using EBS volumes](#): how to use static and dynamic provisioning for jobs that need EBS volumes.

[Using Amazon FSx for Lustre volumes](#): how to use static and dynamic provisioning for jobs that need Amazon FSx for Luster volumes.

[Using Instance store volumes](#): how to use instance store volumes for job processing.

Metastore integration

[Using Hive metastore](#): offers different ways to use Hive metastore.

[Using AWS Glue](#): offers different ways to configure AWS Glue catalog.

Debugging

[Using Spark debugging](#): how to change the log level.

[Connecting to Spark UI on the driver pod](#).

[How to use self-hosted Spark history server with Amazon EMR on EKS](#).

Troubleshooting Amazon EMR on EKS issues

[Troubleshooting](#).

Node placement

[Using Kubernetes node selectors](#) for single-az and other use cases.

[Using Fargate node placement](#).

Performance

[Using Dynamic Resource Allocation \(DRA\)](#).

By default, `spark.dynamicAllocation.preallocateExecutors` is enabled in Amazon EMR Spark. When `spark.dynamicAllocation.initialExecutors` and `spark.dynamicAllocation.minExecutors` are not set, Spark may request a large number of executors at startup based on estimated task counts, even for small workloads. To avoid excessive container churn, use one of the following approaches:

- Set `spark.dynamicAllocation.initialExecutors` or `spark.dynamicAllocation.minExecutors` to a value appropriate for your workload size.
- Set `spark.dynamicAllocation.preallocateExecutors.maxEstimatedTasks` to a lower value to limit the number of executors requested at startup.

- Set `spark.dynamicAllocation.preallocateExecutors` to `false` to disable executor preallocation entirely.

[EKS best practices](#) for the Amazon VPC Container Network Interface plugin (CNI), Cluster Autoscaler, and Core DNS.

Cost optimization

[Using spot instances](#): Amazon EC2 spot instance best practices and how to use the Spark node decommission feature.

Using AWS Outposts

[Running Amazon EMR on EKS using AWS Outposts](#)

Customizing Docker images for Amazon EMR on EKS

You can use customized Docker images with Amazon EMR on EKS. Customizing the Amazon EMR on EKS runtime image provides the following benefits:

- Package application dependencies and runtime environment into a single immutable container that promotes portability and simplifies dependency management for each workload.
- Install and configure packages that are optimized to your workloads. These packages may not be widely available in the public distribution of Amazon EMR runtimes.
- Integrate Amazon EMR on EKS with current established build, test, and deployment processes within your organization, including local development and testing.
- Apply established security processes, such as image scanning, that meet compliance and governance requirements within your organization.

Topics

- [How to customize Docker images](#)
- [Details for selecting a base image URI](#)
- [Considerations for customizing images](#)

How to customize Docker images

Follow these steps to customize Docker images for Amazon EMR on EKS. The steps show you how to get a base image, customize and publish it, and submit a workload using the image.

- [Prerequisites](#)
- [Step 1: Retrieve a base image from Amazon Elastic Container Registry \(Amazon ECR\)](#)
- [Step 2: Customize a base image](#)
- [Step 3: \(Optional but recommended\) Validate a custom image](#)
- [Step 4: Publish a custom image](#)
- [Step 5: Submit a Spark workload in Amazon EMR using a custom image](#)

Note

Other options you may want to consider when customizing Docker images are customizing for interactive endpoints, which you do to ensure you have your required dependencies, or using multi-architectural container images:

- [Customize Docker images for interactive endpoints](#)
- [Work with multi-architecture images](#)

Prerequisites

- Complete the [Setting up Amazon EMR on EKS](#) steps for Amazon EMR on EKS.
- Install Docker in your environment. For more information, see [Get Docker](#).

Step 1: Retrieve a base image from Amazon Elastic Container Registry (Amazon ECR)

The base image contains the Amazon EMR runtime and connectors that are used to access other AWS services. For Amazon EMR 6.9.0 and higher, you can get the base images from the Amazon ECR Public Gallery. Browse the gallery to find the image link and pull the image to your local workspace. For example, for Amazon EMR 7.13.0 release, the following `docker pull` command gets you the latest standard base image. You can replace `emr-7.13.0:latest` with `emr-7.13.0-spark-rapids:latest` to retrieve the image that has Nvidia RAPIDS accelerator. You can also replace `emr-7.13.0:latest` with `emr-7.13.0-java11:latest` to retrieve the image with Java 11 runtime.

```
docker pull public.ecr.aws/emr-on-eks/spark/emr-7.13.0:latest
```

If you would like to retrieve the base image for a Amazon EMR 6.9.0 or earlier releases, or if you prefer to retrieve from Amazon ECR registry accounts in each Region, use the following steps:

1. Choose a base image URI. The image URI follows this format, *ECR-registry-account*.dkr.ecr.*Region*.amazonaws.com/spark/*container-image-tag*, as the following example demonstrates.

```
895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/emr-6.6.0:latest
```

To choose a base image in your Region, see [Details for selecting a base image URI](#).

2. Log in to the Amazon ECR repository where the base image is stored. Replace `895885662937` and `us-west-2` with the Amazon ECR registry account and the AWS Region you have selected.

```
aws ecr get-login-password --region us-west-2 | docker login --username AWS --password-stdin 895885662937.dkr.ecr.us-west-2.amazonaws.com
```

3. Pull the base image into your local Workspace. Replace `emr-6.6.0:latest` with the container image tag you have selected.

```
docker pull 895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/emr-6.6.0:latest
```

Step 2: Customize a base image

Follow these steps to customize the base image you have pulled from Amazon ECR.

1. Create a new Dockerfile on your local Workspace.
2. Edit the Dockerfile you just created and add the following content. This Dockerfile uses the container image you have pulled from `895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/emr-6.6.0:latest`.

```
FROM 895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/emr-6.6.0:latest
USER root
### Add customization commands here ###
USER hadoop:hadoop
```

3. Add commands in the Dockerfile to customize the base image. For example, add a command to install Python libraries, as the following Dockerfile demonstrates.

```
FROM 895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/emr-6.6.0:latest
USER root
RUN pip3 install --upgrade boto3 pandas numpy // For python 3
USER hadoop:hadoop
```

4. From the same directory where the Dockerfile is created, run the following command to build the Docker image. Provide a name for the Docker image, for example, *emr6.6_custom*.

```
docker build -t emr6.6_custom .
```

Step 3: (Optional but recommended) Validate a custom image

We recommend that you test the compatibility of your custom image before publishing it. You can use the [Amazon EMR on EKS custom image CLI](#) to check if your image has the required file structures and correct configurations for running on Amazon EMR on EKS.

Note

The Amazon EMR on EKS custom image CLI cannot confirm that your image is free of error. Use caution when removing dependencies from the base images.

Take the following steps to validate your custom image.

1. Download and install Amazon EMR on EKS custom image CLI. For more information, see [Amazon EMR on EKS custom image CLI Installation Guide](#).
2. Run the following command to test the installation.

```
emr-on-eks-custom-image --version
```

The following shows an example of the output.

```
Amazon EMR on EKS Custom Image CLI  
Version: x.xx
```

3. Run the following command to validate your custom image.

```
emr-on-eks-custom-image validate-image -i image_name -r release_version [-  
t image_type]
```

- *-i* specifies the local image URI that needs to be validated. This can be the image URI, any name or tag that you defined for your image.

- `-r` specifies the exact release version for the base image, for example, `emr-6.6.0-latest`.
- `-t` specifies the image type. If this is a Spark image, input `spark`. The default value is `spark`. The current Amazon EMR on EKS custom image CLI version only supports Spark runtime images.

If you run the command successfully and the custom image meets all the required configurations and file structures, the returned output displays the results of all of the tests, as the following example demonstrates.

```
Amazon EMR on EKS Custom Image Test
Version: x.xx
... Checking if docker cli is installed
... Checking Image Manifest
[INFO] Image ID: xxx
[INFO] Created On: 2021-05-17T20:50:07.986662904Z
[INFO] Default User Set to hadoop:hadoop : PASS
[INFO] Working Directory Set to /home/hadoop : PASS
[INFO] Entrypoint Set to /usr/bin/entrypoint.sh : PASS
[INFO] SPARK_HOME is set with value: /usr/lib/spark : PASS
[INFO] JAVA_HOME is set with value: /etc/alternatives/jre : PASS
[INFO] File Structure Test for spark-jars in /usr/lib/spark/jars: PASS
[INFO] File Structure Test for hadoop-files in /usr/lib/hadoop: PASS
[INFO] File Structure Test for hadoop-jars in /usr/lib/hadoop/lib: PASS
[INFO] File Structure Test for bin-files in /usr/bin: PASS
... Start Running Sample Spark Job
[INFO] Sample Spark Job Test with local:///usr/lib/spark/examples/jars/spark-
examples.jar : PASS
-----
Overall Custom Image Validation Succeeded.
-----
```

If the custom image doesn't meet the required configurations or file structures, error messages occur. The returned output provides information about the incorrect configurations or file structures.

Step 4: Publish a custom image

Publish the new Docker image to your Amazon ECR registry.

1. Run the following command to create an Amazon ECR repository for storing your Docker image. Provide a name for your repository, for example, `emr6.6_custom_repo`. Replace `us-west-2` with your Region.

```
aws ecr create-repository \  
  --repository-name emr6.6_custom_repo \  
  --image-scanning-configuration scanOnPush=true \  
  --region us-west-2
```

For more information, see [Create a repository](#) in the *Amazon ECR User Guide*.

2. Run the following command to authenticate to your default registry.

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

For more information, see [Authenticate to your default registry](#) in the *Amazon ECR User Guide*.

3. Tag and publish an image to the Amazon ECR repository you created.

Tag the image.

```
docker tag emr6.6_custom aws_account_id.dkr.ecr.us-  
west-2.amazonaws.com/emr6.6_custom_repo
```

Push the image.

```
docker push aws_account_id.dkr.ecr.us-west-2.amazonaws.com/emr6.6_custom_repo
```

For more information, see [Push an image to Amazon ECR](#) in the *Amazon ECR User Guide*.

Step 5: Submit a Spark workload in Amazon EMR using a custom image

After a custom image is built and published, you can submit an Amazon EMR on EKS job using a custom image.

First, create a `start-job-run-request.json` file and specify the `spark.kubernetes.container.image` parameter to reference the custom image, as the following example JSON file demonstrates.

Note

You can use `local://` scheme to refer to files available in the custom image as shown with `entryPoint` argument in the JSON snippet below. You can also use the `local://` scheme to refer to application dependencies. All files and dependencies that are referred using `local://` scheme must already be present at the specified path in the custom image.

```
{
  "name": "spark-custom-image",
  "virtualClusterId": "virtual-cluster-id",
  "executionRoleArn": "execution-role-arn",
  "releaseLabel": "emr-6.6.0-latest",
  "jobDriver": {
    "sparkSubmitJobDriver": {
      "entryPoint": "local:///usr/lib/spark/examples/jars/spark-examples.jar",
      "entryPointArguments": [
        "10"
      ],
      "sparkSubmitParameters": "--class org.apache.spark.examples.SparkPi --conf spark.kubernetes.container.image=123456789012.dkr.ecr.us-west-2.amazonaws.com/emr6.6_custom_repo"
    }
  }
}
```

You can also reference the custom image with `applicationConfiguration` properties as the following example demonstrates.

```
{
  "name": "spark-custom-image",
  "virtualClusterId": "virtual-cluster-id",
  "executionRoleArn": "execution-role-arn",
  "releaseLabel": "emr-6.6.0-latest",
  "jobDriver": {
    "sparkSubmitJobDriver": {
      "entryPoint": "local:///usr/lib/spark/examples/jars/spark-examples.jar",
```

```
    "entryPointArguments": [
      "10"
    ],
    "sparkSubmitParameters": "--class org.apache.spark.examples.SparkPi"
  },
  "configurationOverrides": {
    "applicationConfiguration": [
      {
        "classification": "spark-defaults",
        "properties": {
          "spark.kubernetes.container.image": "123456789012.dkr.ecr.us-  
west-2.amazonaws.com/emr6.6_custom_repo"
        }
      }
    ]
  }
}
```

Then run the `start-job-run` command to submit the job.

```
aws emr-containers start-job-run --cli-input-json file:///./start-job-run-request.json
```

In the JSON examples above, replace *emr-6.6.0-latest* with your Amazon EMR release version. We strongly recommend that you use the `-latest` release version to ensure that the selected version contains the latest security updates. For more information about Amazon EMR release versions and their image tags, see [Details for selecting a base image URI](#).

Note

You can use `spark.kubernetes.driver.container.image` and `spark.kubernetes.executor.container.image` to specify a different image for driver and executor pods.

Customize Docker images for interactive endpoints

You can also customize Docker images for interactive endpoints so that you can run customized base kernel images. This helps you ensure that you have the dependencies you need when you run interactive workloads from EMR Studio.

1. Follow the [Steps 1-4](#) outlined above to customize a Docker image. For Amazon EMR 6.9.0 releases and later, you can get the base image URI from Amazon ECR Public Gallery. For releases before Amazon EMR 6.9.0, you can get the image in Amazon ECR Registry accounts in each AWS Region, and the only difference is the base image URI in your Dockerfile. The base image URI follows the format:

```
ECR-registry-account.dkr.ecr.Region.amazonaws.com/notebook-spark/container-image-tag
```

You need to use notebook-spark in the base image URI, instead of spark. The base image contains the Spark runtime and the notebook kernels that run with it. For more information about selecting Regions and container image tags, see [Details for selecting a base image URI](#).

Note

Currently only overrides of base images are supported and introducing completely new kernels of other types than the base images AWS provides is not supported.

2. Create an interactive endpoint that can be used with the custom image.

First, create a JSON file called `custom-image-managed-endpoint.json` with the following contents.

```
{
  "name": "endpoint-name",
  "virtualClusterId": "virtual-cluster-id",
  "type": "JUPYTER_ENTERPRISE_GATEWAY",
  "releaseLabel": "emr-6.6.0-latest",
  "executionRoleArn": "execution-role-arn",
  "certificateArn": "certificate-arn",
  "configurationOverrides": {
    "applicationConfiguration": [
      {
        "classification": "jupyter-kernel-overrides",
        "configurations": [
          {
            "classification": "python3",
            "properties": {
              "container-image": "123456789012.dkr.ecr.us-west-2.amazonaws.com/custom-notebook-python:latest"
            }
          }
        ]
      }
    ]
  }
}
```



```

        }
      },
      {
        "classification": "spark-python-kubernetes",
        "properties": {
          "container-image": "123456789012.dkr.ecr.us-
west-2.amazonaws.com/custom-notebook-spark:latest"
        }
      }
    ]
  }
}

```

Next, create an interactive endpoint using the configurations specified in the JSON file, as the following example demonstrates.

```
aws emr-containers create-managed-endpoint --cli-input-json custom-image-managed-
endpoint.json
```

For more information, see [Create an interactive endpoint for your virtual cluster](#).

3. Connect to the interactive endpoint via EMR Studio. For more information, see [Connecting from Studio](#).

Work with multi-architecture images

Amazon EMR on EKS supports multi-architecture container images for Amazon Elastic Container Registry (Amazon ECR). For more information, see [Introducing multi-architecture container images for Amazon ECR](#).

Amazon EMR on EKS custom images support both AWS Graviton-based EC2 instances and non-Graviton-based EC2 instances. The Graviton-based images are stored in the same image repositories in Amazon ECR as non-Graviton-based images.

For example, to inspect the Docker manifest list for 6.6.0 images, run the following command.

```
docker manifest inspect 895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/
emr-6.6.0:latest
```

Here is the output. The arm64 architecture is for Graviton instance. The amd64 is for non-Graviton instance.

```
{
  "schemaVersion": 2,
  "mediaType": "application/vnd.docker.distribution.manifest.list.v2+json",
  "manifests": [
    {
      "mediaType": "application/vnd.docker.distribution.manifest.v2+json",
      "size": 1805,
      "digest":
"xxx123:6b971cb47d11011ab3d45fff925e9442914b4977ae0f9fbcdf5cfa99a7593f0",
      "platform": {
        "architecture": "arm64",
        "os": "linux"
      }
    },
    {
      "mediaType": "application/vnd.docker.distribution.manifest.v2+json",
      "size": 1805,
      "digest":
"xxx123:6f2375582c9c57fa9838c1d3a626f1b4fc281e287d2963a72dfe0bd81117e52f",
      "platform": {
        "architecture": "amd64",
        "os": "linux"
      }
    }
  ]
}
```

Take the following steps to create multi-architecture images:

1. Create a Dockerfile with the following contents so that you can pull the arm64 image.

```
FROM --platform=arm64 895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/
emr-6.6.0:latest
USER root

RUN pip3 install boto3 // install customizations here
USER hadoop:hadoop
```

2. Follow the instructions at [Introducing multi-architecture container images for Amazon ECR](#) to build a multi-architecture image.

Note

You must create arm64 images on arm64 instances. Similarly, you must build amd64 images on amd64 instances.

You can also build multi-architecture images without building on each specific instance type with the Docker buildx command. For more information, see [Leverage multi-CPU architecture support](#).

3. After you build the multi-architecture image, you can submit a job with the same `spark.kubernetes.container.image` parameter and point it to the image. In a heterogeneous cluster with both AWS Graviton-based and non-Graviton-based EC2 instances, the instance determines the correct architecture image based on the instance architecture that pulls the image.

Details for selecting a base image URI

Note

For Amazon EMR 6.9.0 releases and later, you can retrieve the base image from Amazon ECR Public Gallery, so you don't need to construct the base image URI as the instructions on this page direct. To find the container image tag for your base image, refer to the [release notes page](#) for the corresponding release of Amazon EMR on EKS.

The base Docker images that you can select are stored in Amazon Elastic Container Registry (Amazon ECR). The image URI follows this format: *ECR-registry-account.dkr.ecr.Region.amazonaws.com/spark/container-image-tag*, as the following example demonstrates.

```
895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/emr-7.13.0:latest
```

The image URI for interactive endpoints follows this format: *ECR-registry-account.dkr.ecr.Region.amazonaws.com/notebook-spark/container-image-tag*, as the following example demonstrates. You need to use `notebook-spark` in the base image URI, instead of `spark`.

```
895885662937.dkr.ecr.us-west-2.amazonaws.com/notebook-spark/emr-7.13.0:latest
```

Similarly, for non-Spark python3 images for interactive endpoints, the image URI is *ECR-registry-account.dkr.ecr.Region.amazonaws.com/notebook-python/container-image-tag*. The following example URI is correctly formatted:

```
895885662937.dkr.ecr.us-west-2.amazonaws.com/notebook-python/emr-7.13.0:latest
```

To find the container image tag for your base image, refer to the [release notes page](#) for the corresponding release of Amazon EMR on EKS.

Amazon ECR registry accounts by Region

To avoid high network latency, pull a base image from your closest AWS Region. Select the Amazon ECR registry account that corresponds with the Region that you pull the image from based on the following table.

Regions	Amazon ECR registry accounts	
ap-east-1	736135916053	
ap-northeast-1	059004520145	
ap-northeast-2	996579266876	
ap-northeast-3	705689932349	
ap-southeast-3	946962994502	
ap-south-1	235914868574	
ap-south-2	691480105545	
ap-southeast-1	671219180197	
ap-southeast-2	038297999601	
ca-central-1	351826393999	

Regions	Amazon ECR registry accounts	
eu-central-1	107292555468	
eu-central-2	314408114945	
eu-north-1	830386416364	
eu-west-1	483788554619	
eu-west-2	118780647275	
eu-west-3	307523725174	
eu-south-1	238014973495	
eu-south-2	350796622945	
il-central-1	395734710648	
me-south-1	008085056818	
me-central-1	818935616732	
sa-east-1	052806832358	
us-gov-west-1	299385240661	
us-gov-east-1	299393998622	
us-east-1	755674844232	
us-east-2	711395599931	
us-west-1	608033475327	
us-west-2	895885662937	
af-south-1	358491847878	
cn-north-1	068337069695	

Regions	Amazon ECR registry accounts	
cn-northwest-1	068420816659	

Considerations for customizing images

When you customize Docker images, you can choose the exact runtime for your job at a granular level. Consider these best practices when you use this feature. These include considerations for security, configuration, and mounting an image:

- Security is a shared responsibility between AWS and you. You're responsible for security patching the binaries that you add to the image. Follow the [Amazon EMR on EKS security best practices](#), especially [Get the latest security updates for custom images](#) and [Apply principle of least privilege](#).
- When you customize a base image, you must change the Docker user to `hadoop:hadoop` so that the jobs do not run with the root user.
- Amazon EMR on EKS mounts files on top of the configurations for the image, such as the `spark-defaults.conf`, at run time. To override these configuration files, we recommend that you use the `applicationOverrides` parameter during the job submission and not directly modify the files in the custom image.
- Amazon EMR on EKS mounts certain folders at run time. Any modifications that you make to these folders aren't available in the container. If you want to add an application or its dependencies for custom images, we recommend that you choose a directory that isn't part of the following predefined paths:
 - `/var/log/fluentd`
 - `/var/log/spark/user`
 - `/var/log/spark/apps`
 - `/mnt`
 - `/tmp`
 - `/home/hadoop`
- You can upload your customized image to any Docker-compatible repository, such as Amazon ECR, Docker Hub, or a private enterprise repository. For more information on how to configure

the Amazon EKS cluster authentication with the selected Docker repository, see [Pull an Image from a Private Registry](#).

 Important

Standard Amazon EMR on EKS pricing applies whenever pods use an image derived from the Amazon EMR runtime, even if you store the customized image in your own private registry and use your own orchestration method (for example, the StartJobRun API, the Spark operator, or spark-submit). For more information, see [Amazon EMR on EKS pricing](#).

Running Flink jobs with Amazon EMR on EKS

Amazon EMR releases 6.13.0 and higher support Amazon EMR on EKS with Apache Flink, or the Flink Kubernetes operator, as a job submission model for Amazon EMR on EKS. With Amazon EMR on EKS with Apache Flink, you can deploy and manage Flink applications with the Amazon EMR release runtime on your own Amazon EKS clusters. Once you deploy the Flink Kubernetes operator in your Amazon EKS cluster, you can directly submit Flink applications with the operator. The operator manages the lifecycle of Flink applications.

Topics

- [Setting up and using the Flink Kubernetes operator](#)
- [Using Flink Native Kubernetes](#)
- [Customizing Docker images for Flink and FluentD](#)
- [Monitoring Flink Kubernetes operator and Flink jobs](#)
- [How Flink supports high availability and job resiliency](#)
- [Using Autoscaler for Flink applications](#)
- [Maintenance and troubleshooting for Flink jobs on Amazon EMR on EKS](#)
- [Supported releases for Amazon EMR on EKS with Apache Flink](#)

Setting up and using the Flink Kubernetes operator

The following pages describe how to set up and use the Flink Kubernetes operator to run Flink jobs with Amazon EMR on EKS. The topics available include required prerequisites, how to set up your environment, and running a Flink application on Amazon EMR on EKS.

Topics

- [Setting up the Flink Kubernetes operator for Amazon EMR on EKS](#)
- [Installing the Flink Kubernetes operator for Amazon EMR on EKS](#)
- [Run a Flink application](#)
- [Security role permissions for running a Flink application](#)
- [Uninstalling the Flink Kubernetes operator for Amazon EMR on EKS](#)

Setting up the Flink Kubernetes operator for Amazon EMR on EKS

Complete the following tasks to get set up before you install the Flink Kubernetes operator on Amazon EKS. If you've already signed up for Amazon Web Services (AWS) and have used Amazon EKS, you are almost ready to use Amazon EMR on EKS. Complete the following tasks to get set up for the Flink operator on Amazon EKS. If you've already completed any of the prerequisites, you can skip those and move on to the next one.

- [Install or update to the latest version of the AWS CLI](#) – If you've already installed the AWS CLI, confirm that you have the latest version.
- [Set up kubectl and eksctl](#) – eksctl is a command line tool that you use to communicate with Amazon EKS.
- [Install Helm](#) – The Helm package manager for Kubernetes helps you install and manage applications on your Kubernetes cluster.
- [Get started with Amazon EKS – eksctl](#) – Follow the steps to create a new Kubernetes cluster with nodes in Amazon EKS.
- [Choose an Amazon EMR release label \(release 6.13.0 or higher\)](#) – the Flink Kubernetes operator is supported with Amazon EMR releases 6.13.0 and higher.
- [Enable IAM Roles for Service Accounts \(IRSA\) on the Amazon EKS cluster.](#)
- [Create a job execution role.](#)
- [Update the trust policy of the job execution role .](#)
- Create an operator execution role. This step is optional. You can use the same role for Flink jobs and operator. If you want to have a different IAM role for your operator, you can create a separate role.
- Update the trust policy of the operator execution role. You must explicitly add one trust policy entry for the roles you want to use for the Amazon EMR Flink Kubernetes operator service account. You can follow this example format:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Federated": "arn:aws:iam::AWS_ACCOUNT_ID:oidc-provider/OIDC_PROVIDER"
      }
    }
  ]
}
```

```

    },
    "Action": "sts:AssumeRoleWithWebIdentity",
    "Condition": {
      "StringEquals": {
        "OIDC_PROVIDER:aud": "sts.amazonaws.com",
        "OIDC_PROVIDER:sub": "system:serviceaccount:emr:emr-containers-sa-
flink-operator"
      }
    }
  }
]
}

```

Installing the Flink Kubernetes operator for Amazon EMR on EKS

This topic helps you start to use the Flink Kubernetes operator on Amazon EKS by preparing a Flink deployment.

Install the Kubernetes operator

Use the following steps to install the Kubernetes operator for Apache Flink.

1. If you haven't already, complete the steps in [the section called "Setting up"](#).
2. Install the *cert-manager* (once per Amazon EKS cluster) to enable adding the webhook component.

```
kubectl apply -f https://github.com/cert-manager/cert-manager/releases/download/
v1.12.0/cert-manager.yaml
```

3. Install the Helm chart.

```

export VERSION=7.13.0 # The Amazon EMR release version
export NAMESPACE=The Kubernetes namespace to deploy the operator

helm install flink-kubernetes-operator \
oci://public.ecr.aws/emr-on-eks/flink-kubernetes-operator \
--version $VERSION \
--namespace $NAMESPACE

```

Example output:

```
NAME: flink-kubernetes-operator
LAST DEPLOYED: Tue May 31 17:38:56 2022
NAMESPACE: $NAMESPACE
STATUS: deployed
REVISION: 1
TEST SUITE: None
```

- Wait for the deployment to be complete and verify the chart installation.

```
kubectl wait deployment flink-kubernetes-operator --namespace $NAMESPACE --for
condition=Available=True --timeout=30s
```

- You should see the following message when deployment is complete.

```
deployment.apps/flink-kubernetes-operator condition met
```

- Use the following command to see the deployed operator.

```
helm list --namespace $NAMESPACE
```

The following shows example output, where the app version `x.y.z-amzn-n` would correspond with the Flink operator version for your Amazon EMR on EKS release. For more information, see [Supported releases for Amazon EMR on EKS with Apache Flink](#).

NAME	STATUS	CHART	NAMESPACE	REVISION	UPDATED	APP VERSION
flink-kubernetes-operator-0500	EST	deployed	\$NAMESPACE	1	2023-02-22 16:43:45.24148	x.y.z-amzn-n

Upgrade the Kubernetes operator

To upgrade the version of the Flink operator, follow these steps:

- Uninstall the old `flink-kubernetes-operator`: `helm uninstall flink-kubernetes-operator -n <NAMESPACE>`.
- Delete CRD (since helm will not automatically delete the old CRD): `kubectl delete crd flinkdeployments.flink.apache.org flinksessionjobs.flink.apache.org`.

3. Re-install flink-kubernetes-operator with the newer version.

Run a Flink application

With Amazon EMR 6.13.0 and higher, you can run a Flink application with the Flink Kubernetes operator in Application mode on Amazon EMR on EKS. With Amazon EMR 6.15.0 and higher, you can also run a Flink application in Session mode. This page describes both methods that you can use to run a Flink application with Amazon EMR on EKS.

Note

You must have an Amazon S3 bucket created to store the high-availability metadata when you submit your Flink job. If you don't want to use this feature, you can disable it. It's enabled by default.

Prerequisite – Before you can run a Flink application with the Flink Kubernetes operator, complete the steps in [the section called “Setting up”](#) and [the section called “Install the Kubernetes operator”](#).

Application mode

With Amazon EMR 6.13.0 and higher, you can run a Flink application with the Flink Kubernetes operator in Application mode on Amazon EMR on EKS.

1. Create a FlinkDeployment definition file `basic-example-app-cluster.yaml` like in the following example. If you activated and use one of the [opt-in AWS Regions](#), make sure you uncomment and configure the configuration `fs.s3a.endpoint.region`.

```
apiVersion: flink.apache.org/v1beta1
kind: FlinkDeployment
metadata:
  name: basic-example-app-cluster
spec:
  flinkConfiguration:
    taskmanager.numberOfTaskSlots: "2"
    #fs.s3a.endpoint.region: OPT_IN_AWS_REGION_NAME
    state.checkpoints.dir: CHECKPOINT_S3_STORAGE_PATH
    state.savepoints.dir: SAVEPOINT_S3_STORAGE_PATH
  flinkVersion: v1_17
  executionRoleArn: JOB_EXECUTION_ROLE_ARN
```

```

emrReleaseLabel: "emr-6.13.0-flink-latest" # 6.13 or higher
jobManager:
  storageDir: HIGH_AVAILABILITY_STORAGE_PATH
  resource:
    memory: "2048m"
    cpu: 1
taskManager:
  resource:
    memory: "2048m"
    cpu: 1
job:
  # if you have your job jar in S3 bucket you can use that path as well
  jarURI: local:///opt/flink/examples/streaming/StateMachineExample.jar
  parallelism: 2
  upgradeMode: savepoint
  savepointTriggerNonce: 0
monitoringConfiguration:
  cloudWatchMonitoringConfiguration:
    logGroupName: LOG_GROUP_NAME

```

2. Submit the Flink deployment with the following command. This will also create a FlinkDeployment object named basic-example-app-cluster.

```
kubectl create -f basic-example-app-cluster.yaml -n <NAMESPACE>
```

3. Access the Flink UI.

```
kubectl port-forward deployments/basic-example-app-cluster 8081 -n NAMESPACE
```

4. Open localhost:8081 to view your Flink jobs locally.
5. Clean up the job. Remember to clean up the S3 artifacts that were created for this job, such as checkpointing, high-availability, savepointing metadata, and CloudWatch logs.

For more information on submitting applications to Flink through the Flink Kubernetes operator, see [Flink Kubernetes operator examples](#) in the apache/flink-kubernetes-operator folder on GitHub.

Session mode

With Amazon EMR 6.15.0 and higher, you can run a Flink application with the Flink Kubernetes operator in Session mode on Amazon EMR on EKS.

1. Create a FlinkDeployment definition file named `basic-example-app-cluster.yaml` like in the following example. If you activated and use one of the [opt-in AWS Regions](#), make sure you uncomment and configure the configuration `fs.s3a.endpoint.region`.

```

apiVersion: flink.apache.org/v1beta1
kind: FlinkDeployment
metadata:
  name: basic-example-session-cluster
spec:
  flinkConfiguration:
    taskmanager.numberOfTaskSlots: "2"
    #fs.s3a.endpoint.region: OPT_IN_AWS_REGION_NAME
    state.checkpoints.dir: CHECKPOINT_S3_STORAGE_PATH
    state.savepoints.dir: SAVEPOINT_S3_STORAGE_PATH
  flinkVersion: v1_17
  executionRoleArn: JOB_EXECUTION_ROLE_ARN
  emrReleaseLabel: "emr-6.15.0-flink-latest"
  jobManager:
    storageDir: HIGH_AVAILABILITY_S3_STORAGE_PATH
    resource:
      memory: "2048m"
      cpu: 1
  taskManager:
    resource:
      memory: "2048m"
      cpu: 1
  monitoringConfiguration:
    s3MonitoringConfiguration:
      logUri:
    cloudWatchMonitoringConfiguration:
      logGroupName: LOG_GROUP_NAME

```

2. Submit the Flink deployment with the following command. This will also create a FlinkDeployment object named `basic-example-session-cluster`.

```
kubectl create -f basic-example-app-cluster.yaml -n NAMESPACE
```

3. Use the following command to confirm that the session cluster LIFECYCLE is STABLE:

```
kubectl get flinkdeployments.flink.apache.org basic-example-session-cluster -n NAMESPACE
```

The output should be similar to the following example:

NAME	JOB STATUS	LIFECYCLE STATE
basic-example-session-cluster		STABLE

4. Create a FlinkSessionJob custom definition resource file `basic-session-job.yaml` with the following example content:

```
apiVersion: flink.apache.org/v1beta1
kind: FlinkSessionJob
metadata:
  name: basic-session-job
spec:
  deploymentName: basic-session-deployment
  job:
    # If you have your job jar in an S3 bucket you can use that path.
    # To use jar in S3 bucket, set
    # OPERATOR_EXECUTION_ROLE_ARN (--set emrContainers.operatorExecutionRoleArn=
    $OPERATOR_EXECUTION_ROLE_ARN)
    # when you install Spark operator
    jarURI: https://repo1.maven.org/maven2/org/apache/flink/flink-examples-
    streaming_2.12/1.16.1/flink-examples-streaming_2.12-1.16.1-TopSpeedWindowing.jar
    parallelism: 2
    upgradeMode: stateless
```

5. Submit the Flink session job with the following command. This will create a FlinkSessionJob object `basic-session-job`.

```
kubectl apply -f basic-session-job.yaml -n $NAMESPACE
```

6. Use the following command to confirm that the session cluster LIFECYCLE is STABLE, and the JOB STATUS is RUNNING:

```
kubectl get flinkdeployments.flink.apache.org basic-example-session-cluster -
n NAMESPACE
```

The output should be similar to the following example:

NAME	JOB STATUS	LIFECYCLE STATE
basic-example-session-cluster	RUNNING	STABLE

7. Access the Flink UI.

```
kubectl port-forward deployments/basic-example-session-cluster 8081 -n NAMESPACE
```

8. Open `localhost:8081` to view your Flink jobs locally.
9. Clean up the job. Remember to clean up the S3 artifacts that were created for this job, such as checkpointing, high-availability, savepointing metadata, and CloudWatch logs.

Security role permissions for running a Flink application

This topic describes security roles for deploying and running a Flink application. There are two roles required to manage a deployment and to create and manage jobs, the operator role and job role. This topic introduces them and lists their permissions.

Role based access control

To deploy the operator and run Flink jobs, we must create two Kubernetes roles: one operator and one job role. Amazon EMR creates the two roles by default when you install the operator.

Operator role

We use the operator role to manage `flinkdeployments` to create and manage the JobManager for each Flink job and other resources, like services.

The operator role's default name is `emr-containers-sa-flink-operator` and requires the following permissions.

```
rules:
- apiGroups:
  - ""
  resources:
  - pods
  - services
  - events
  - configmaps
  - secrets
  - serviceaccounts
  verbs:
  - '*'
- apiGroups:
  - rbac.authorization.k8s.io
```



```
resources:
- roles
- rolebindings
verbs:
- '*'
- apiGroups:
  - apps
  resources:
  - deployments
  - deployments/finalizers
  - replicaset
  verbs:
  - '*'
- apiGroups:
  - extensions
  resources:
  - deployments
  - ingresses
  verbs:
  - '*'
- apiGroups:
  - flink.apache.org
  resources:
  - flinkdeployments
  - flinkdeployments/status
  - flinksessionjobs
  - flinksessionjobs/status
  verbs:
  - '*'
- apiGroups:
  - networking.k8s.io
  resources:
  - ingresses
  verbs:
  - '*'
- apiGroups:
  - coordination.k8s.io
  resources:
  - leases
  verbs:
  - '*'
```

Job role

The JobManager uses the job role to create and manage TaskManagers and ConfigMaps for each job.

```
rules:
- apiGroups:
  - ""
  resources:
  - pods
  - configmaps
  verbs:
  - '*'
- apiGroups:
  - apps
  resources:
  - deployments
  - deployments/finalizers
  verbs:
  - '*'
```

Uninstalling the Flink Kubernetes operator for Amazon EMR on EKS

Follow these steps to uninstall the Flink Kubernetes operator.

1. Delete the operator.

```
helm uninstall flink-kubernetes-operator -n <NAMESPACE>
```

2. Delete Kubernetes resources that Helm doesn't uninstall.

```
kubectl delete serviceaccounts, roles, rolebindings -l emr-
containers.amazonaws.com/component=flink.operator --namespace <namespace>
kubectl delete crd flinkdeployments.flink.apache.org
flinksessionjobs.flink.apache.org
```

3. (Optional) Delete the cert-manager.

```
kubectl delete -f https://github.com/jetstack/cert-manager/releases/download/
v1.12.0/cert-manager.yaml
```

Using Flink Native Kubernetes

Amazon EMR releases 6.13.0 and higher support Flink Native Kubernetes as a command-line tool that you can use to submit and execute Flink applications to an Amazon EMR on EKS cluster.

Topics

- [Setting up Flink Native Kubernetes for Amazon EMR on EKS](#)
- [Getting started with Flink native Kubernetes for Amazon EMR on EKS](#)
- [Flink JobManager service account security requirements for Native Kubernetes](#)

Setting up Flink Native Kubernetes for Amazon EMR on EKS

Complete the following tasks to get set up before you can run an application with the Flink CLI on Amazon EMR on EKS. If you've already signed up for Amazon Web Services (AWS) and have used Amazon EKS, you are almost ready to use Amazon EMR on EKS. If you've already completed any of the prerequisites, you can skip those and move on to the next one.

- [Install or update to the latest version of the AWS CLI](#) – If you've already installed the AWS CLI, confirm that you have the latest version.
- [Get started with Amazon EKS – eksctl](#) – Follow the steps to create a new Kubernetes cluster with nodes in Amazon EKS.
- [Select an Amazon EMR base image URI \(release 6.13.0 or higher\)](#) – the Flink Kubernetes command is supported with Amazon EMR releases 6.13.0 and higher.
- Confirm that the JobManager service account has appropriate permissions to create and watch TaskManager pods. For more information, see [Flink JobManager service account security requirements for Native Kubernetes](#).
- Set up your local [AWS credentials profile](#).
- [Create or updating a kubeconfig file for an Amazon EKS cluster](#) on which you want to run the Flink applications.

Getting started with Flink native Kubernetes for Amazon EMR on EKS

These steps show you how to configure, set up a service account for, and run a Flink application. Flink Native Kubernetes is used to deploy Flink on a running Kubernetes cluster.

Configure and run a Flink application

Amazon EMR 6.13.0 and higher supports Flink Native Kubernetes for running Flink applications on an Amazon EKS cluster. To run a Flink application, follow these steps:

1. Before you can run a Flink application with the Flink Native Kubernetes command, complete the steps in [the section called "Setting up"](#).
2. [Download and install Flink](#).
3. Set the values for the following environment variables.

```
#Export the FLINK_HOME environment variable to your local installation of Flink
export FLINK_HOME=/usr/local/bin/flink #Will vary depending on your installation
export NAMESPACE=flink
export CLUSTER_ID=flink-application-cluster
export IMAGE=<123456789012.dkr.ecr.sample-AWS Region-.amazonaws.com/flink/
emr-6.13.0-flink:latest>
export FLINK_SERVICE_ACCOUNT=emr-containers-sa-flink
export FLINK_CLUSTER_ROLE_BINDING=emr-containers-crb-flink
```

4. Create a service account to manage Kubernetes resources.

```
kubectl create serviceaccount $FLINK_SERVICE_ACCOUNT -n $NAMESPACE
kubectl create clusterrolebinding $FLINK_CLUSTER_ROLE_BINDING --clusterrole=edit --
serviceaccount=$NAMESPACE:$FLINK_SERVICE_ACCOUNT
```

5. Run the run-application CLI command.

```
$FLINK_HOME/bin/flink run-application \
  --target kubernetes-application \
  -Dkubernetes.namespace=$NAMESPACE \
  -Dkubernetes.cluster-id=$CLUSTER_ID \
  -Dkubernetes.container.image.ref=$IMAGE \
  -Dkubernetes.service-account=$FLINK_SERVICE_ACCOUNT \
  local:///opt/flink/examples/streaming/Iteration.jar
2022-12-29 21:13:06,947 INFO  org.apache.flink.kubernetes.utils.KubernetesUtils
    [] - Kubernetes deployment requires a fixed port. Configuration
    blob.server.port will be set to 6124
2022-12-29 21:13:06,948 INFO  org.apache.flink.kubernetes.utils.KubernetesUtils
    [] - Kubernetes deployment requires a fixed port. Configuration
    taskmanager.rpc.port will be set to 6122
2022-12-29 21:13:07,861 WARN
    org.apache.flink.kubernetes.KubernetesClusterDescriptor    [] - Please note that
```

Flink client operations(e.g. cancel, list, stop, savepoint, etc.) won't work from outside the Kubernetes cluster since 'kubernetes.rest-service.exposed.type' has been set to ClusterIP.

2022-12-29 21:13:07,868 INFO

```
org.apache.flink.kubernetes.KubernetesClusterDescriptor [] - Create flink
application cluster flink-application-cluster successfully, JobManager Web
Interface: http://flink-application-cluster-rest.flink:8081
```

6. Examine the created Kubernetes resources.

```
kubectl get all -n <namespace>
NAME READY STATUS RESTARTS AGE
pod/flink-application-cluster-546687cb47-w2p2z 1/1 Running 0 3m37s
pod/flink-application-cluster-taskmanager-1-1 1/1 Running 0 3m24s

NAME TYPE CLUSTER-IP EXTERNAL-IP PORT(S) AGE
service/flink-application-cluster ClusterIP None <none> 6123/TCP,6124/TCP 3m38s
service/flink-application-cluster-rest ClusterIP 10.100.132.158 <none> 8081/TCP
3m38s


NAME READY UP-TO-DATE AVAILABLE AGE
deployment.apps/flink-application-cluster 1/1 1 1 3m38s

NAME DESIRED CURRENT READY AGE
replicaset.apps/flink-application-cluster-546687cb47 1 1 1 3m38s
```

7. Port forward to 8081.

```
kubectl port-forward service/flink-application-cluster-rest 8081 -n <namespace>
Forwarding from 127.0.0.1:8081 -> 8081
```

8. Locally access the Flink UI.



9. Delete the Flink application.

```
kubectl delete deployment.apps/flink-application-cluster -n <namespace>
deployment.apps "flink-application-cluster" deleted
```

For more information about submitting applications to Flink, see [Native Kubernetes](#) in the Apache Flink documentation.

Flink JobManager service account security requirements for Native Kubernetes

The Flink JobManager pod uses a Kubernetes service account to access the Kubernetes API server to create and watch TaskManager pods. The JobManager service account must have appropriate permissions to create/delete TaskManager pods and allow the TaskManager to watch leader ConfigMaps to retrieve the address of JobManager and ResourceManager in your cluster.

The following rules apply to this service account.

```
rules:
- apiGroups:
  - ""
  resources:
  - pods
  verbs:
  - "*"
- apiGroups:
```

```
- ""
resources:
- services
verbs:
- "*"
- apiGroups:
- ""
resources:
- configmaps
verbs:
- "*"
- apiGroups:
- "apps"
resources:
- deployments
verbs:
- "*"

```

Customizing Docker images for Flink and FluentD

Take the following steps to customize Docker images for Amazon EMR on EKS with Apache Flink or FluentD images. These include technical guidance for getting a base image, customizing it, publishing it, and submitting a workload.

Topics

- [Prerequisites](#)
- [Step 1: Retrieve a base image from Amazon Elastic Container Registry](#)
- [Step 2: Customize a base image](#)
- [Step 3: Publish your custom image](#)
- [Step 4: Submit a Flink workload in Amazon EMR using a custom image](#)

Prerequisites

Before you customize your Docker image, make sure that you have completed the following prerequisites:

- Completed the [Setting up the Flink Kubernetes operator for Amazon EMR on EKS](#) steps.
- Installed Docker in your environment. For more information, see [Get Docker](#).

Step 1: Retrieve a base image from Amazon Elastic Container Registry

The base image contains the Amazon EMR runtime and connectors that you need to access other AWS services. If you're using Amazon EMR on EKS with Flink version 6.14.0 or higher, you can get the base images from the Amazon ECR Public Gallery. Browse the gallery to find the image link and pull the image to your local workspace. For example, for the Amazon EMR 6.14.0 release, the following `docker pull` command returns the latest standard base image. Replace `emr-6.14.0:latest` with the release version you want.

```
docker pull public.ecr.aws/emr-on-eks/flink/emr-6.14.0-flink:latest
```

The following are links to the Flink gallery image and Fluentd gallery image:

- [emr-on-eks/flink/emr-6.14.0-flink](#)
- [emr-on-eks/fluentd/emr-6.14.0-fluentd](#)

Step 2: Customize a base image

The following steps describe how to customize the base image you pulled from Amazon ECR.

1. Create a new Dockerfile on your local Workspace.
2. Edit the Dockerfile and add the following content. This Dockerfile uses the container image you pulled from `public.ecr.aws/emr-on-eks/flink/emr-7.13.0-flink:latest`.

```
FROM public.ecr.aws/emr-on-eks/flink/emr-7.13.0-flink:latest
USER root
### Add customization commands here ###
USER hadoop:hadoop
```

Use the following configuration if you're using Fluentd.

```
FROM public.ecr.aws/emr-on-eks/fluentd/emr-7.13.0:latest
USER root
### Add customization commands here ###
USER hadoop:hadoop
```


3. Add commands in the Dockerfile to customize the base image. The following command demonstrates how to install Python libraries.

```
FROM public.ecr.aws/emr-on-eks/flink/emr-7.13.0-flink:latest
USER root
RUN pip3 install --upgrade boto3 pandas numpy // For python 3
USER hadoop:hadoop
```

4. In the same directory of where you created DockerFile, run the following command to build the Docker image. The field you supply following the -t flag is your custom name for the image.

```
docker build -t <YOUR_ACCOUNT_ID>.dkr.ecr.<YOUR_ECR_REGION>.amazonaws.com/
<ECR_REPO>:<ECR_TAG>
```

Step 3: Publish your custom image

You can now publish the new Docker image to your Amazon ECR registry.

1. Run the following command to create an Amazon ECR repository to store your Docker image. Provide a name for your repository, such as `emr_custom_repo`. For more information, see [Create a repository](#) in the Amazon Elastic Container Registry User Guide.

```
aws ecr create-repository \
    --repository-name emr_custom_repo \
    --image-scanning-configuration scanOnPush=true \
    --region <AWS_REGION>
```

2. Run the following command to authenticate to your default registry. For more information, see [Authenticate to your default registry](#) in the Amazon Elastic Container Registry User Guide.

```
aws ecr get-login-password --region <AWS_REGION> | docker login --username AWS --
password-stdin <AWS_ACCOUNT_ID>.dkr.ecr.<YOUR_ECR_REGION>.amazonaws.com
```

3. Push the image. For more information, see [Push an image to Amazon ECR](#) in the Amazon Elastic Container Registry User Guide.

```
docker push <YOUR_ACCOUNT_ID>.dkr.ecr.<YOUR_ECR_REGION>.amazonaws.com/
<ECR_REPO>:<ECR_TAG>
```

Step 4: Submit a Flink workload in Amazon EMR using a custom image

Make the following changes to your FlinkDeployment spec to use a custom image. To do so, enter your own image in the spec.image line of your deployment spec.

```
apiVersion: flink.apache.org/v1beta1
kind: FlinkDeployment
metadata:
  name: basic-example
spec:
  flinkVersion: v1_18
  image: <YOUR_ACCOUNT_ID>.dkr.ecr.<YOUR_ECR_REGION>.amazonaws.com/
  <ECR_REPO>:<ECR_TAG>
  imagePullPolicy: Always
  flinkConfiguration:
    taskmanager.numberOfTaskSlots: "1"
```

To use a custom image for your Fluentd job, enter your own image in the monitoringConfiguration.image line of your deployment spec.

```
monitoringConfiguration:
  image: <YOUR_ACCOUNT_ID>.dkr.ecr.<YOUR_ECR_REGION>.amazonaws.com/
  <ECR_REPO>:<ECR_TAG>
  cloudWatchMonitoringConfiguration:
    logGroupName: flink-log-group
    logStreamNamePrefix: custom-fluentd
```

Monitoring Flink Kubernetes operator and Flink jobs

This section describes several ways that you can monitor your Flink jobs with Amazon EMR on EKS. These include integrating Flink with the Amazon Managed Service for Prometheus, using the *Flink Web Dashboard*, which provides job status and metrics, or using a monitoring configuration to send log data to Amazon S3 and Amazon CloudWatch.

Topics

- [Use Amazon Managed Service for Prometheus to monitor Flink jobs](#)
- [Use the Flink UI to monitor Flink jobs](#)
- [Use monitoring configuration to monitor Flink Kubernetes operator and Flink jobs](#)

Use Amazon Managed Service for Prometheus to monitor Flink jobs

You can integrate Apache Flink with Amazon Managed Service for Prometheus (management portal). Amazon Managed Service for Prometheus supports ingesting metrics from Amazon Managed Service for Prometheus servers in clusters running on Amazon EKS. Amazon Managed Service for Prometheus works together with a Prometheus server already running on your Amazon EKS cluster. Running Amazon Managed Service for Prometheus integration with Amazon EMR Flink operator will automatically deploy and configure a Prometheus server to integrate with Amazon Managed Service for Prometheus.

1. [Create an Amazon Managed Service for Prometheus Workspace](#). This workspace serves as an ingestion endpoint. You will need the remote write URL later.
2. Set up IAM roles for service accounts.

For this method of onboarding, use IAM roles for the service accounts in the Amazon EKS cluster where the Prometheus server is running. These roles are also called *service roles*.

If you don't already have the roles, [set up service roles for the ingestion of metrics from Amazon EKS clusters](#).

Before you continue, create an IAM role called `amp-iamproxy-ingest-role`.

3. Install the Amazon EMR Flink Operator with Amazon Managed Service for Prometheus.

Now that you have an Amazon Managed Service for Prometheus workspace, a dedicated IAM role for Amazon Managed Service for Prometheus, and the necessary permissions, you can install the Amazon EMR Flink operator.

Create an `enable-amp.yaml` file. This file lets you use a custom configuration to override Amazon Managed Service for Prometheus settings. Make sure to use your own roles.

```
kube-prometheus-stack:
  prometheus:
    serviceAccount:
      create: true
      name: "amp-iamproxy-ingest-service-account"
      annotations:
        eks.amazonaws.com/role-arn: "arn:aws:iam::<AWS_ACCOUNT_ID>:role/amp-iamproxy-ingest-role"
    remoteWrite:
```

```
- url: <AMAZON_MANAGED_PROMETHEUS_REMOTE_WRITE_URL>
sigv4:
  region: <AWS_REGION>
queueConfig:
  maxSamplesPerSend: 1000
  maxShards: 200
  capacity: 2500
```

Use the `Helm Install --set` command to pass overrides to the `flink-kubernetes-operator` chart.

```
helm upgrade -n <namespace> flink-kubernetes-operator \
  oci://public.ecr.aws/emr-on-eks/flink-kubernetes-operator \
  --set prometheus.enabled=true
-f enable-amp.yaml
```

This command automatically installs a Prometheus reporter in the operator on port 9999. Any future `FlinkDeployment` also exposes a metrics port on 9249.

- Flink operator metrics appear in Prometheus under the label `flink_k8soperator_`.
- Flink Task Manager metrics appear in Prometheus under the label `flink_taskmanager_`.
- Flink Job Manager metrics appear in Prometheus under the label `flink_jobmanager_`.

Use the Flink UI to monitor Flink jobs

To monitor the health and performance of a running Flink application, use the *Flink Web Dashboard*. This dashboard provides information about the status of the job, the number of TaskManagers, and the metrics and logs for the job. It also lets you view and modify the configuration of the Flink job, and to interact with the Flink cluster to submit or cancel jobs.

To access the Flink Web Dashboard for a running Flink application on Kubernetes:

1. Use the `kubectl port-forward` command to forward a local port to the port on which the Flink Web Dashboard is running in the Flink application's TaskManager pods. By default, this port is 8081. Replace *deployment-name* with the name of the Flink application deployment from above.

```
kubectl get deployments -n namespace
```

Example output:

```
kubectl get deployments -n flink-namespace
```


NAME	READY	UP-TO-DATE	AVAILABLE	AGE
basic-example	1/1	1	1	11m
flink-kubernetes-operator	1/1	1	1	21h

```
kubectl port-forward deployments/deployment-name 8081 -n namespace
```

- If you want to use a different port locally, use the *local-port*:8081 parameter.

```
kubectl port-forward -n flink deployments/basic-example 8080:8081
```

- In a web browser, navigate to <http://localhost:8081> (or <http://localhost:local-port> if you used a custom local port) to access the Flink Web Dashboard. This dashboard shows information about the running Flink application, such as the status of the job, the number of TaskManagers, and the metrics and logs for the job.



Use monitoring configuration to monitor Flink Kubernetes operator and Flink jobs

Monitoring configuration lets you easily set up log archiving of your Flink application and operator logs to S3 and/or CloudWatch (you can choose either one or both). Doing so adds a FluentD sidecar to your JobManager and TaskManager pods and subsequently forwards these components' logs to your configured sinks.

Note

You must set up IAM Roles for the service account for your Flink operator and your Flink job (Service Accounts) to be able to use this feature, as it requires interacting with other AWS services. You must set this up using IRSA in [Setting up the Flink Kubernetes operator for Amazon EMR on EKS](#).

Flink application logs

You can define this configuration in the following way.

```
apiVersion: flink.apache.org/v1beta1
kind: FlinkDeployment
metadata:
  name: basic-example
spec:
  image: FLINK IMAGE TAG
  imagePullPolicy: Always
  flinkVersion: v1_17
  flinkConfiguration:
    taskmanager.numberOfTaskSlots: "2"
  executionRoleArn: JOB EXECUTION ROLE
  jobManager:
    resource:
      memory: "2048m"
      cpu: 1
  taskManager:
    resource:
      memory: "2048m"
      cpu: 1
  job:
```

```

jarURI: local:///opt/flink/examples/streaming/StateMachineExample.jar
monitoringConfiguration:
  s3MonitoringConfiguration:
    logUri: S3 BUCKET
  cloudWatchMonitoringConfiguration:
    logGroupName: LOG GROUP NAME
    logStreamNamePrefix: LOG GROUP STREAM PREFIX
  sidecarResources:
    limits:
      cpuLimit: 500m
      memoryLimit: 250Mi
  containerLogRotationConfiguration:
    rotationSize: 2GB
    maxFilesToKeep: 10

```

The following are configuration options.

- `s3MonitoringConfiguration` – configuration key to set up forwarding to S3
- `logUri` (required) – the S3 bucket path of where you want to store your logs.
- The path on S3 once the logs are uploaded will look like the following.
- No log rotation enabled:

```
s3://${logUri}/${POD NAME}/STDOUT or STDERR.gz
```

- Log rotation is enabled. You can use both a rotated file and a current file (one without the date stamp).

```
s3://${logUri}/${POD NAME}/STDOUT or STDERR.gz
```

The following format is an incrementing number.

```
s3://${logUri}/${POD NAME}/stdout_YYYYMMDD_index.gz
```

- The following IAM permissions are required to use this forwarder.

```

{
  "Effect": "Allow",
  "Action": [
    "s3:PutObject"
  ],
  "Resource": [

```

```

        "S3_BUCKET_URI/*",
        "S3_BUCKET_URI"
    ]
}

```

- `cloudWatchMonitoringConfiguration` – configuration key to set up forwarding to CloudWatch.
- `logGroupName` (required) – name of the CloudWatch log group that you want to send logs to (automatically creates the group if it doesn't exist).
- `logStreamNamePrefix` (optional) – name of the log stream that you want to send logs into. Default value is an empty string. The format is as follows:

```

${logStreamNamePrefix}/${POD_NAME}/STDOUT or STDERR

```

- The following IAM permissions are required to use this forwarder.

```

{
  "Effect": "Allow",
  "Action": [
    "logs:CreateLogStream",
    "logs:CreateLogGroup",
    "logs:PutLogEvents"
  ],
  "Resource": [
    "arn:aws:logs:REGION:ACCOUNT-ID:log-group:{YOUR_LOG_GROUP_NAME}:*",
    "arn:aws:logs:REGION:ACCOUNT-ID:log-group:{YOUR_LOG_GROUP_NAME}"
  ]
}

```

- `sidecarResources` (optional) – the configuration key to set resource limits on the launched Fluentbit sidecar container.
 - `memoryLimit` (optional) – the default value is 512Mi. Adjust according to your needs.
 - `cpuLimit` (optional) – this option doesn't have a default. Adjust according to your needs.
- `containerLogRotationConfiguration` (optional) – controls the container log rotation behavior. It is enabled by default.
 - `rotationSize` (required) – specifies the file size for the log rotation. The range of possible values is from 2KB to 2GB. The numeric unit portion of the `rotationSize` parameter is passed as an integer. Since decimal values aren't supported, you can specify a rotation size of 1.5GB, for example, with the value 1500MB. The default is 2GB.

- `maxFilesToKeep` (required) – specifies the maximum number of files to retain in container after rotation has taken place. The minimum value is 1, and the maximum value is 50. The default is 10.

Flink operator logs

We can also enable log archiving for the operator by using the following options in the values.yaml file in your helm chart installation. You can enable S3, CloudWatch, or both.

```
monitoringConfiguration:
  s3MonitoringConfiguration:
    logUri: "S3-BUCKET"
    totalFileSize: "1G"
    uploadTimeout: "1m"
  cloudWatchMonitoringConfiguration:
    logGroupName: "flink-log-group"
    logStreamNamePrefix: "example-job-prefix-test-2"
  sidecarResources:
    limits:
      cpuLimit: 1
      memoryLimit: 800Mi
    memoryBufferLimit: 700M
```

The following are the available configuration options under `monitoringConfiguration`.

- `s3MonitoringConfiguration` – set this option to archive to S3.
- `logUri` (required) – The S3 bucket path where you want to store your logs.
- The following are formats of what the S3 bucket paths might look like once the logs are uploaded.
- No log rotation enabled.

```
s3://{logUri}/{POD NAME}/OPERATOR or WEBHOOK/STDOUT or STDERR.gz
```

- Log rotation is enabled. You can use both a rotated file and a current file (one without the date stamp).

```
s3://{logUri}/{POD NAME}/OPERATOR or WEBHOOK/STDOUT or STDERR.gz
```

The following format index is an incrementing number.

```
s3://${logUri}/${POD_NAME}/OPERATOR or WEBHOOK/stdout_YYYYMMDD_index.gz
```

- `cloudWatchMonitoringConfiguration` – the configuration key to set up forwarding to CloudWatch.
- `logGroupName` (required) – name of the CloudWatch log group that you want to send logs to. The group automatically gets created if it doesn't exist.
- `logStreamNamePrefix` (optional) – name of the log stream that you want to send logs into. The default value is an empty string. The format in CloudWatch is as follows:

```
${logStreamNamePrefix}/${POD_NAME}/STDOUT or STDERR
```

- `sideCarResources` (optional) – the configuration key to set resource limits on the launched Fluentbit sidecar container.
 - `memoryLimit` (optional) – the memory limit. Adjust according to your needs. The default is 512Mi.
 - `cpuLimit` – the CPU limit. Adjust according to your needs. No default value.
- `containerLogRotationConfiguration` (optional): – controls the container log rotation behavior. It is enabled by default.
 - `rotationSize` (required) – specifies file size for the log rotation. The range of possible values is from 2KB to 2GB. The numeric unit portion of the `rotationSize` parameter is passed as an integer. Since decimal values aren't supported, you can specify a rotation size of 1.5GB, for example, with the value 1500MB. The default is 2GB.
 - `maxFilesToKeep` (required) – specifies the maximum number of files to retain in container after rotation has taken place. The minimum value is 1, and the maximum value is 50. The default is 10.

How Flink supports high availability and job resiliency

The following sections outline how Flink makes jobs more reliable and highly available. It does this through built-in capabilities like Flink high availability and various recovery capabilities if failures occur.

Topics

- [Using high availability \(HA\) for Flink Operators and Flink Applications](#)

- [Optimizing Flink job restart times for task recovery and scaling operations with Amazon EMR on EKS](#)
- [Graceful decommission of Spot Instances with Flink on Amazon EMR on EKS](#)

Using high availability (HA) for Flink Operators and Flink Applications

This topic shows how to configure high availability and describes how it works for a few different use cases. These include when you're using the Job manager and when you're using Flink native kubernetes.

Flink operator high-availability

We enable *high availability* for the Flink Operator so that we can fail-over to a standby Flink Operator to minimize downtime in the operator control loop if failures occur. High availability is enabled by default and the default number of starting operator replicas is 2. You can configure the replicas field in your `values.yaml` file for the helm chart.

The following fields are customizable:

- `replicas` (optional, default is 2): Setting this number to greater than 1 creates other standby Operators and allows for faster recovery of your job.
- `highAvailabilityEnabled` (optional, default is true): Controls whether you want to enable HA. Specifying this parameter as true enables multi AZ deployment support, as well as sets the correct `flink-conf.yaml` parameters.

You can disable HA for your operator by setting the following configuration in your `values.yaml` file.

```
...
imagePullSecrets: []

replicas: 1

# set this to false if you don't want HA
highAvailabilityEnabled: false
...
```

Multi AZ deployment

We create the operator pods in multiple Availability Zones. This is a soft constraint, and your operator pods will be scheduled in the same AZ if you don't have enough resources in a different AZ.

Determining the leader replica

If HA is enabled, the replicas use a lease to determine which of the JMs is the leader and uses a K8s Lease for leader election. You can describe the Lease and look at the `.Spec.Holder Identity` field to determine the current leader

```
kubectl describe lease <Helm Install Release Name>-<NAMESPACE>-lease -n <NAMESPACE> |  
grep "Holder Identity"
```

Flink-S3 Interaction

Configuring access credentials

Please make sure that you have configured IRSA with appropriate IAM permissions to access the S3 bucket.

Fetching job jars from S3 Application mode

The Flink operator also supports fetching applications jars from S3. You just provide the S3 location for the jarURI in your FlinkDeployment specification.

You can also use this feature to download other artifacts like PyFlink scripts. The resulting Python script is dropped under the path `/opt/flink/usr/lib/`.

The following example demonstrates how to use this feature for a PyFlink job. Note the jarURI and args fields.

```
apiVersion: flink.apache.org/v1beta1  
kind: FlinkDeployment  
metadata:  
  name: python-example  
spec:  
  image: <YOUR CUSTOM PYFLINK IMAGE>  
  emrReleaseLabel: "emr-6.12.0-flink-latest"  
  flinkVersion: v1_16  
  flinkConfiguration:  
    taskmanager.numberOfTaskSlots: "1"
```

```

serviceAccount: flink
jobManager:
  highAvailabilityEnabled: false
  replicas: 1
  resource:
    memory: "2048m"
    cpu: 1
taskManager:
  resource:
    memory: "2048m"
    cpu: 1
job:
  jarURI: "s3://<S3-BUCKET>/scripts/pyflink.py" # Note, this will trigger the
artifact download process
  entryClass: "org.apache.flink.client.python.PythonDriver"
  args: ["-pyclientexec", "/usr/local/bin/python3", "-py", "/opt/flink/usrlib/
pyflink.py"]
  parallelism: 1
  upgradeMode: stateless

```

Flink S3 Connectors

Flink comes packaged with two S3 connectors (listed below). The following sections discuss when to use which connector.

Checkpointing: Presto S3 connector

- Set S3 scheme to s3p://
- The recommended connector to use to checkpoint to s3. For more information, see [S3-specific](#) in the Apache Flink documentation.

Example FlinkDeployment specification:

```

apiVersion: flink.apache.org/v1beta1
kind: FlinkDeployment
metadata:
  name: basic-example
spec:
  flinkConfiguration:
    taskmanager.numberOfTaskSlots: "2"
    state.checkpoints.dir: s3p://<BUCKET-NAME>/flink-checkpoint/

```

Reading and writing to S3: Hadoop S3 connector

- Set S3 scheme to `s3://` or `( s3a:// )`
- The recommended connector for reading and writing files from S3 (only S3 connector that implements the [Flinks Filesystem interface](#)).
- By default, we set `fs.s3a.aws.credentials.provider` in the `flink-conf.yaml` file, which is `com.amazonaws.auth.WebIdentityTokenCredentialsProvider`. If you override the default `flink-conf` completely and you are interacting with S3, make sure to use this provider.

Example FlinkDeployment spec

```
apiVersion: flink.apache.org/v1beta1
kind: FlinkDeployment
metadata:
  name: basic-example
spec:
  job:
    jarURI: local:///opt/flink/examples/streaming/WordCount.jar
    args: [ "--input", "s3a://<INPUT_BUCKET>/PATH", "--output", "s3a://<OUTPUT_BUCKET>/PATH" ]
    parallelism: 2
    upgradeMode: stateless
```

Flink Job Manager

High Availability (HA) for Flink Deployments allow jobs to continue making progress even if a transient error is encountered and your JobManager crashes. The jobs will restart but from the last successful checkpoint with HA enabled. Without HA enabled, Kubernetes will restart your JobManager, but your job will start as a fresh job and will lose its progress. After configuring HA, we can tell Kubernetes to store the HA metadata in a persistent storage to reference in case of a transient failure in the JobManager and then resume our jobs from the last successful checkpoint.

HA is enabled by default for your Flink jobs (the replica count is set to 2, which will require you to provide an S3 storage location for HA metadata to persist).

HA configs

```
apiVersion: flink.apache.org/v1beta1
kind: FlinkDeployment
metadata:
```

```
name: basic-example
spec:
  flinkConfiguration:
    taskmanager.numberOfTaskSlots: "2"
    executionRoleArn: "<JOB EXECUTION ROLE ARN>"
    emrReleaseLabel: "emr-6.13.0-flink-latest"
  jobManager:
    resource:
      memory: "2048m"
      cpu: 1
    replicas: 2
    highAvailabilityEnabled: true
    storageDir: "s3://<S3 PERSISTENT STORAGE DIR>"
  taskManager:
    resource:
      memory: "2048m"
      cpu: 1
```

The following are descriptions for the above HA configs in Job Manager (defined under `.spec.jobManager`):

- `highAvailabilityEnabled` (optional, default is true): Set this to `false` if you don't want HA enabled and don't want to use the provided HA configurations. You can still manipulate the `"replicas"` field to manually configure HA.
- `replicas` (optional, default is 2): Setting this number to greater than 1 creates other standby JobManagers and allows for faster recovery of your job. If you disable HA, you must set replica count to 1, or you will keep getting validation errors (only 1 replica is supported if HA is not enabled).
- `storageDir` (required): Because we use replica count as 2 by default, we have to provide a persistent storageDir. Currently this field only accepts S3 paths as the storage location.

Pod locality

If you enable HA, we also try to colocate Job Manager pods and Task Manager pods in the same AZ, which leads to improved performance (reduced network latency by having pods in same AZs). This is a best-effort process, meaning if you don't have enough resources in the AZ where the majority of your Job Manager and Task Manager pods are scheduled, the remaining pods will still be scheduled but might end up on a node outside of this AZ.

Determining the leader replica

If HA is enabled, the replicas use a lease to determine which of the JMs is the leader and uses a K8s Configmap as the datastore to store this metadata. If you want to determine the leader, you can look at the content of the Configmap and look at the key `org.apache.flink.k8s.leader.restserver` under `data` to find the K8s pod with the IP address. You can also use the following bash commands.

```
ip=$(kubectl get configmap -n <NAMESPACE> <JOB-NAME>-cluster-config-map -o json | jq -r ".data[\"org.apache.flink.k8s.leader.restserver\"]" | awk -F: '{print $2}' | awk -F '/' '{print $3}')
kubectl get pods -n NAMESPACE -o json | jq -r ".items[] | select(.status.podIP == \"$ip\") | .metadata.name"
```

Flink job - native Kubernetes

Amazon EMR 6.13.0 and higher supports Flink native Kubernetes for running Flink applications in high-availability mode on an Amazon EKS cluster.

Note

You must have an Amazon S3 bucket created to store the high-availability metadata when you submit your Flink job. If you don't want to use this feature, you can disable it. It's enabled by default.

To turn on the Flink high-availability feature, provide the following Flink parameters when you [run the run-application CLI command](#). The parameters are defined below the example.

```
-Dhigh-availability.type=kubernetes \
-Dhigh-availability.storageDir=S3://DOC-EXAMPLE-STORAGE-BUCKET \
-
Dfs.s3a.aws.credentials.provider="com.amazonaws.auth.WebIdentityTokenCredentialsProvider" \
-Dkubernetes.jobmanager.replicas=3 \
-Dkubernetes.cluster-id=example-cluster
```

- **Dhigh-availability.storageDir** – The Amazon S3 bucket where you want to store the high-availability metadata for your job.

Dkubernetes.jobmanager.replicas – The number of Job Manager pods to create as an integer greater than 1.

`Dkubernetes.cluster-id` – A unique ID that identifies the Flink cluster.

Optimizing Flink job restart times for task recovery and scaling operations with Amazon EMR on EKS

When a task fails or when a scaling operation occurs, Flink attempts to re-execute the task from the last completed checkpoint. The restart process could take a minute or longer to execute, depending on the size of the checkpoint state and the number of parallel tasks. During the restart period, backlog tasks can accumulate for the job. There are some ways though, that Flink optimizes the speed of recovery and restart of execution graphs to improve job stability.

This page describes some of the ways that Amazon EMR Flink can improve the job restart time during task recovery or scaling operations on spot instances. Spot instances are unused compute capacity that's available at a discount. It has unique behaviors, including occasional interruptions, so it's important to understand how Amazon EMR on EKS handles these, including how Amazon EMR on EKS carries out decommissioning and job restarts.

Topics

- [Task-local recovery](#)
- [Task-local recovery by Amazon EBS volume mount](#)
- [Generic log-based incremental checkpoint](#)
- [Fine-grained recovery](#)
- [Combined restart mechanism in adaptive scheduler](#)

Task-local recovery

Note

Task-local recovery is supported with Flink on Amazon EMR on EKS 6.14.0 and higher.

With Flink checkpoints, each task produces a snapshot of its state that Flink writes to distributed storage like Amazon S3. In cases of recovery, the tasks restore their state from the distributed storage. Distributed storage provides fault tolerance and can redistribute the state during rescaling because it's accessible to all nodes.

However, a remote distributed store also has a disadvantage: all tasks must read their state from a remote location over the network. This can result in long recovery times for large states during task recovery or scaling operations.

This problem of long recovery time is solved by *task-local recovery*. Tasks write their state on checkpoint into a secondary storage that is local to the task, such as on a local disk. They also store their state in the primary storage, or Amazon S3 in our case. During recovery, the scheduler schedules the tasks on the same Task Manager where the tasks ran earlier so that they can recover from the local state store instead of reading from the remote state store. For more information, see [Task-Local Recovery](#) in the *Apache Flink Documentation*.

Our benchmark tests with sample jobs have shown that the recovery time has been reduced from minutes to a few seconds with task-local recovery enabled.

To enable task-local recovery, set the following configurations in your `flink-conf.yaml` file. Specify the checkpointing interval value in milliseconds.

```
state.backend.local-recovery: true
state.backend: hasmap or rocksdb
state.checkpoints.dir: s3://STORAGE-BUCKET-PATH/checkpoint
execution.checkpointing.interval: 15000
```

Task-local recovery by Amazon EBS volume mount

Note

Task-local recovery by Amazon EBS is supported with Flink on Amazon EMR on EKS 6.15.0 and higher.

With Flink on Amazon EMR on EKS, you can automatically provision Amazon EBS volumes to the TaskManager pods for task local recovery. The default overlay mount comes with 10 GB volume, which is sufficient for jobs with a lower state. Jobs with large states can enable the *automatic EBS volume mount* option. The TaskManager pods are automatically created and mounted during pod creation and removed during pod deletion.

Use the following steps to enable automatic EBS volume mount for Flink in Amazon EMR on EKS:

1. Export the values for the following variables that you'll use in upcoming steps.

```
export AWS_REGION=aa-example-1
export FLINK_EKS_CLUSTER_NAME=my-cluster
export AWS_ACCOUNT_ID=111122223333
```

2. Create or update a kubeconfig YAML file for your cluster.

```
aws eks update-kubeconfig --name $FLINK_EKS_CLUSTER_NAME --region $AWS_REGION
```

3. Create an IAM service account for the Amazon EBS Container Storage Interface (CSI) driver on your Amazon EKS cluster.

```
eksctl create iamserviceaccount \
  --name ebs-csi-controller-sa \
  --namespace kube-system \
  --region $AWS_REGION \
  --cluster $FLINK_EKS_CLUSTER_NAME \
  --role-name TLR_${AWS_REGION}_${FLINK_EKS_CLUSTER_NAME} \
  --role-only \
  --attach-policy-arn arn:aws:iam::aws:policy/service-role/
AmazonEBSCSIDriverPolicy \
  --approve
```

4. Create the Amazon EBS CSI driver with the following command:

```
eksctl create addon \
  --name aws-ebs-csi-driver \
  --region $AWS_REGION \
  --cluster $FLINK_EKS_CLUSTER_NAME \
  --service-account-role-arn arn:aws:iam::${AWS_ACCOUNT_ID}:role/TLR_
${AWS_REGION}_${FLINK_EKS_CLUSTER_NAME}
```

5. Create the Amazon EBS storage class with the following command:

```
cat # EOF # storage-class.yaml
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: ebs-sc
provisioner: ebs.csi.aws.com
volumeBindingMode: WaitForFirstConsumer
EOF
```

And then apply the class:

```
kubectl apply -f storage-class.yaml
```

6. Helm install the Amazon EMR Flink Kubernetes operator with options to create a service account. This creates the `emr-containers-sa-flink` to use in the Flink deployment.

```
helm install flink-kubernetes-operator flink-kubernetes-operator/ \
  --set jobServiceAccount.create=true \
  --set rbac.jobRole.create=true \
  --set rbac.jobRoleBinding.create=true
```

7. To submit the Flink job and enable the automatic provision of EBS volumes for task-local recovery, set the following configurations in your `flink-conf.yaml` file. Adjust the size limit for the state size of the job. Set `serviceAccount` to `emr-containers-sa-flink`. Specify the checkpointing interval value in milliseconds. And omit the `executionRoleArn`.

```
flinkConfiguration:
  task.local-recovery.ebs.enable: true
  kubernetes.taskmanager.local-recovery.persistentVolumeClaim.sizeLimit: 10Gi
  state.checkpoints.dir: s3://BUCKET-PATH/checkpoint
  state.backend.local-recovery: true
  state.backend: hasmap or rocksdb
  state.backend.incremental: "true"
  execution.checkpointing.interval: 15000
  serviceAccount: emr-containers-sa-flink
```

When you're ready to delete the Amazon EBS CSI driver plugin, use the following commands:

```
# Detach Attached Policy
aws iam detach-role-policy --role-name TLR_${AWS_REGION}_${FLINK_EKS_CLUSTER_NAME}
--policy-arn arn:aws:iam::aws:policy/service-role/AmazonEBSCSIDriverPolicy
# Delete the created Role
aws iam delete-role --role-name TLR_${AWS_REGION}_${FLINK_EKS_CLUSTER_NAME}
# Delete the created service account
eksctl delete iamserviceaccount --name ebs-csi-controller-sa --namespace kube-system
--cluster $FLINK_EKS_CLUSTER_NAME --region $AWS_REGION
# Delete Addon
eksctl delete addon --name aws-ebs-csi-driver --cluster $FLINK_EKS_CLUSTER_NAME --
region $AWS_REGION
```

```
# Delete the EBS storage class
kubectl delete -f storage-class.yaml
```

Generic log-based incremental checkpoint

Note

Generic log-based incremental checkpointing is supported with Flink on Amazon EMR on EKS 6.14.0 and higher.

Generic log-based incremental checkpointing was added in Flink 1.16 to improve the speed of checkpoints. A faster checkpoint interval often results in a reduction of recovery work because fewer events need to be reprocessed after recovery. For more information, see [Improving speed and stability of checkpointing with generic log-based incremental checkpoints](#) on the *Apache Flink Blog*.

With sample jobs, our benchmark tests have shown that the checkpoint time reduced from minutes to a few seconds with the generic log-based incremental checkpoint.

To enable generic log-based incremental checkpoints, set the following configurations in your `flink-conf.yaml` file. Specify the checkpointing interval value in milliseconds.

```
state.backend.changelog.enabled: true
state.backend.changelog.storage: filesystem
dstl.dfs.base-path: s3://bucket-path/changelog
state.backend.local-recovery: true
state.backend: rocksdb
state.checkpoints.dir: s3://bucket-path/checkpoint
execution.checkpointing.interval: 15000
```

Fine-grained recovery

Note

Fine-grained recovery support for the default scheduler is supported with Flink on Amazon EMR on EKS 6.14.0 and higher. Fine-grained recovery support in the adaptive scheduler is available with Flink on Amazon EMR on EKS 6.15.0 and higher.

When a task fails during execution, Flink resets the entire execution graph and triggers complete re-execution from the last completed checkpoint. This is more expensive than just re-executing the failed tasks. Fine-grained recovery restarts only the pipeline-connected component of the failed task. In the following example, the job graph has 5 vertices (A to E). All connections between the vertices are pipelined with pointwise distribution, and the `parallelism.default` for the job is set to 2.

```
A # B # C # D # E
```

For this example, there are a total of 10 tasks running. The first pipeline (a1 to e1) runs on a TaskManager (TM1), and the second pipeline (a2 to e2) runs on another TaskManager (TM2).

```
a1 # b1 # c1 # d1 # e1  
a2 # b2 # c2 # d2 # e2
```

There are two pipelined connected components: a1 # e1, and a2 # e2. If either TM1 or TM2 fails, the failure impacts only the 5 tasks in the pipeline where the TaskManager was running. The restart strategy only starts the affected pipelined component.

Fine-grained recovery works only with perfectly parallel Flink jobs. It's not supported with `keyBy()` or `redistribute()` operations. For more information, see [FLIP-1: Fine Grained Recovery from Task Failures](#) in the *Flink Improvement Proposal* Jira project.

To enable fine-grained recovery, set the following configurations in your `flink-conf.yaml` file.

```
jobmanager.execution.failover-strategy: region  
restart-strategy: exponential-delay or fixed-delay
```

Combined restart mechanism in adaptive scheduler

Note

The combined restart mechanism in adaptive scheduler is supported with Flink on Amazon EMR on EKS 6.15.0 and higher.

Adaptive scheduler can adjust the parallelism of the job based on available slots. It automatically reduces the parallelism if not enough slots are available to fit the configured job parallelism. If new

slots become available, the job is scaled up again to the configured job parallelism. An adaptive scheduler avoids downtime on the job when there are not enough resources available. This is the supported scheduler for Flink Autoscaler. We recommend adaptive scheduler with Amazon EMR Flink for these reasons. However, adaptive schedulers might do multiple restarts within a short period of time, one restart for every new resource added. This could lead to a performance drop in the job.

With Amazon EMR 6.15.0 and higher, Flink has a combined restart mechanism in adaptive scheduler that opens a restart window when the first resource is added, and then waits until the configured window interval of the default 1 minute. It performs a single restart when there are sufficient resources available to run the job with configured parallelism or when the interval times out.

With sample jobs, our benchmark tests have shown that this feature processes 10% of records more than the default behavior when you use adaptive scheduler and Flink autoscaler.

To enable the combined restart mechanism, set the following configurations in your `flink-conf.yaml` file.

```
jobmanager.adaptive-scheduler.combined-restart.enabled: true
jobmanager.adaptive-scheduler.combined-restart.window-interval: 1m
```

Graceful decommission of Spot Instances with Flink on Amazon EMR on EKS

Flink with Amazon EMR on EKS can improve the job restart time during task recovery or scaling operations.

Overview

Amazon EMR on EKS releases 6.15.0 and higher support graceful decommission of Task Managers on Spot Instances in Amazon EMR on EKS with Apache Flink. As part of this feature, Amazon EMR on EKS with Flink provides the following capabilities:

- **Just-in-time checkpointing** – Flink streaming jobs can respond to Spot Instance interruption, perform just-in-time (JIT) checkpoint of the running jobs, and prevent scheduling of additional tasks on these Spot Instances. JIT checkpoint is supported with default and adaptive scheduler.
- **Combined restart mechanism** – A combined restart mechanism makes a best-effort attempt to restart the job after it reaches target resource parallelism or the end of the current configured

window. This also prevents consecutive job restarts that might be caused by multiple Spot Instance terminations. Combined restart mechanism is available with adaptive scheduler only.

These capabilities provide the following benefits:

- You can leverage Spot Instances to run Task Managers and reduce cluster expenditure.
- Improved liveness for Spot Instance Task Manager results in higher resilience and more efficient job scheduling.
- Your Flink jobs will have more uptime because there will be less restarts from Spot Instance termination.

How graceful decommissioning works

Consider the following example: you provision an Amazon EMR on EKS cluster running Apache Flink, and you specify On-Demand nodes for Job Manager, and Spot Instance nodes for Task Manager. Two minutes before termination, Task Manager receives an interruption notice.

In this scenario, the Job Manager would handle the Spot Instance interruption signal, block scheduling of additional tasks on the Spot Instance, and initiate JIT checkpointing for the streaming job.

Then, the Job Manager would restart the job graph only after there is sufficient availability of new resources to satisfy current job parallelism in the current restart interval window. The restart window interval is decided on the basis of Spot Instance replacement duration, creation of new Task Manager pods, and registration with Job Manager.

Prerequisites

To use graceful decommissioning, create and run a streaming job on an Amazon EMR on EKS cluster running Apache Flink. Enable Adaptive Scheduler and Task Managers scheduled on at least one Spot Instance, as shown in the following example. You should use On-Demand nodes for Job Manager, and you can use On-Demand nodes for Task Managers as long as there's at least one Spot Instance, too.

```
apiVersion: flink.apache.org/v1beta1
kind: FlinkDeployment
metadata:
  name: deployment_name
```



```
spec:
  flinkVersion: v1_17
  flinkConfiguration:
    taskmanager.numberOfTaskSlots: "2"
    cluster.taskmanager.graceful-decommission.enabled: "true"
    execution.checkpointing.interval: "240s"
    jobmanager.adaptive-scheduler.combined-restart.enabled: "true"
    jobmanager.adaptive-scheduler.combined-restart.window-interval : "1m"
  serviceAccount: flink
  jobManager:
    resource:
      memory: "2048m"
      cpu: 1
    nodeSelector:
      'eks.amazonaws.com/capacityType': 'ON_DEMAND'
  taskManager:
    resource:
      memory: "2048m"
      cpu: 1
    nodeSelector:
      'eks.amazonaws.com/capacityType': 'SPOT'
  job:
    jarURI: flink_job_jar_path
```

Configuration

This section covers most of the configurations that you can specify for your decommissioning needs.

Key	Description	Default value	Acceptable values
<code>cluster.taskmanager.graceful-decommission.enabled</code>	Enable graceful decommission of Task Manager.	true	true, false
<code>jobmanager.adaptive</code>	Enable combined restart mechanism in Adaptive Scheduler.	false	true, false

Key	Description	Default value	Acceptable values
e-scheduler.combined-restart.enabled			
jobmanager.adaptive-scheduler.combined-restart.window-interval	The combined restart window interval to perform merged restarts for the job. An integer without a unit is interpreted as milliseconds.	1m	Examples: 30, 60s, 3m, 1h

Using Autoscaler for Flink applications

The operator autoscaler can help ease backpressure by collecting metrics from Flink jobs and automatically adjusting parallelism on a job vertex level. The following is an example of what your configuration might look like:

```
apiVersion: flink.apache.org/v1beta1
kind: FlinkDeployment
metadata:
  ...
spec:
  ...
  flinkVersion: v1_18
  flinkConfiguration:
    job.autoscaler.enabled: "true"
    job.autoscaler.stabilization.interval: 1m
    job.autoscaler.metrics.window: 5m
    job.autoscaler.target.utilization: "0.6"
    job.autoscaler.target.utilization.boundary: "0.2"
    job.autoscaler.restart.time: 2m
    job.autoscaler.catch-up.duration: 5m
    pipeline.max-parallelism: "720"
  ...
```

This configuration uses default values for the latest release of Amazon EMR. If you use other versions, you might have different values.

 Note

As of Amazon EMR 7.2.0, you don't need to include the prefix `kubernetes.operator` in your configuration. If you use 7.1.0 or lower, you must use the prefix before each configuration. For example, you must specify `kubernetes.operator.job.autoscaler.scaling.enabled`.

The following are configuration options for the autoscaler.

- `job.autoscaler.scaling.enabled` – specifies whether to enable vertex scaling execution by the autoscaler. The default is `true`. If you disable this configuration, the autoscaler only collects metrics and evaluates the suggested parallelism for each vertex but doesn't upgrade the jobs.
- `job.autoscaler.stabilization.interval` – the stabilization period in which no new scaling will be executed. Default is 5 minutes.
- `job.autoscaler.metrics.window` – the scaling metrics aggregation window size. The larger the window, the more smooth and stability, but the autoscaler might be slower to react to sudden load changes. Default is 15 minutes. We recommend you experiment by using a value between 3 to 60 minutes.
- `job.autoscaler.target.utilization` – the target vertex utilization to provide stable job performance and some buffer for load fluctuations. The default is `0.7` targeting 70% utilization/load for the job vertexes.
- `job.autoscaler.target.utilization.boundary` – the target vertex utilization boundary that serves as extra buffer to avoid immediate scaling on load fluctuations. Default is `0.3`, which means 30% deviation from the target utilization is allowed before triggering a scaling action.
- `ob.autoscaler.restart.time` – the expected time to restart the application. Default is 5 minutes.
- `job.autoscaler.catch-up.duration` – the expected time to catch up, meaning fully processing any backlog after a scaling operation completes. Default is 5 minutes. By lowering the catch-up duration, the autoscaler has to reserve more extra capacity for the scaling actions.
- `pipeline.max-parallelism` – the maximum parallelism the autoscaler can use. The autoscaler ignores this limit if it is higher than the max parallelism configured in the Flink config or directly on each operator. Default is `-1`. Note that the autoscaler computes the

parallelism as a divisor of the max parallelism number therefore it is recommended to choose max parallelism settings that have a lot of divisors instead of relying on the Flink provided defaults. We recommend using multiples of 60 for this configuration, such as 120, 180, 240, 360, 720 etc.

For a more detailed configuration reference page, see [Autoscaler configuration](#).

Autoscaler parameter autotuning

This section describes auto-tuning behavior for various Amazon EMR versions. It also goes into detail regarding different auto-scaling configurations.

Note

Amazon EMR 7.2.0 and higher uses the open source configuration `job.autoscaler.restart.time-tracking.enabled` to enable **rescale time estimation**. Rescale time estimation has the same functionality as Amazon EMR autotuning, so you don't have to manually assign empirical values to the restart time. You can still use Amazon EMR autotuning if you're using Amazon EMR 7.1.0 or lower.

7.2.0 and higher

Amazon EMR 7.2.0 and higher measures the actual required restart time to apply autoscaling decisions. In releases 7.1.0 and lower, you had to use the configuration `job.autoscaler.restart.time` to manually configure estimated maximum restart time. By using the configuration `job.autoscaler.restart.time-tracking.enabled`, you only need to enter a restart time for the first scaling. Afterwards, the operator records the actual restart time and will use it for subsequent scalings.

To enable this tracking, use the following command:

```
job.autoscaler.restart.time-tracking.enabled: true
```

The following are the related configurations for rescale time estimation.

Configuration	Required	Default	Description
<code>job.autoscaler.restart.time-tracking.enabled</code>	No	False	Indicates whether the Flink Autoscaler should automatically tune configurations over time to optimize scaling decisions. Note that the Autoscaler can only autotune the Autoscaler parameter <code>restart.time</code> .
<code>job.autoscaler.restart.time</code>	No	5m	The expected restart time that Amazon EMR on EKS uses until the operator can determine the actual restart time from previous scalings.
<code>job.autoscaler.restart.time-tracking.limit</code>	No	15m	The maximum observed restart time when <code>job.autoscaler.restart.time-tracking.enabled</code> is set to true.

The following is an example deployment spec you can use to try out rescale time estimation:

```

apiVersion: flink.apache.org/v1beta1
kind: FlinkDeployment
metadata:
  name: autoscaling-example
spec:
  flinkVersion: v1_18
  flinkConfiguration:

    # Autoscaler parameters
    job.autoscaler.enabled: "true"
    job.autoscaler.scaling.enabled: "true"
    job.autoscaler.stabilization.interval: "5s"
    job.autoscaler.metrics.window: "1m"

```

```

job.autoscaler.restart.time-tracking.enabled: "true"
job.autoscaler.restart.time: "2m"
job.autoscaler.restart.time-tracking.limit: "10m"

jobmanager.scheduler: adaptive
taskmanager.numberOfTaskSlots: "1"
pipeline.max-parallelism: "12"

executionRoleArn: <JOB_ARN>
emrReleaseLabel: emr-7.13.0-flink-latest
jobManager:
  highAvailabilityEnabled: false
  storageDir: s3://<s3_bucket>/flink/autoscaling/ha/
  replicas: 1
  resource:
    memory: "1024m"
    cpu: 0.5
taskManager:
  resource:
    memory: "1024m"
    cpu: 0.5
job:
  jarURI: s3://<s3_bucket>/some-job-with-back-pressure
  parallelism: 1
  upgradeMode: stateless

```

To simulate backpressure, use the following deployment spec.

```

job:
  jarURI: s3://<s3_bucket>/pyflink-script.py
  entryClass: "org.apache.flink.client.python.PythonDriver"
  args: ["-py", "/opt/flink/usrlib/pyflink-script.py"]
  parallelism: 1
  upgradeMode: stateless

```

Upload the following Python script to your S3 bucket.

```

import logging
import sys
import time
import random

```

```

from pyflink.datastream import StreamExecutionEnvironment
from pyflink.table import StreamTableEnvironment

TABLE_NAME="orders"
QUERY=f"""
CREATE TABLE {TABLE_NAME} (
    id INT,
    order_time AS CURRENT_TIMESTAMP,
    WATERMARK FOR order_time AS order_time - INTERVAL '5' SECONDS
)
WITH (
    'connector' = 'datagen',
    'rows-per-second'='10',
    'fields.id.kind'='random',
    'fields.id.min'='1',
    'fields.id.max'='100'
);
"""

def create_backpressure(i):
    time.sleep(2)
    return i

def autoscaling_demo():
    env = StreamExecutionEnvironment.get_execution_environment()
    t_env = StreamTableEnvironment.create(env)
    t_env.execute_sql(QUERY)
    res_table = t_env.from_path(TABLE_NAME)

    stream = t_env.to_data_stream(res_table) \
        .shuffle().map(lambda x: create_backpressure(x)) \
        .print()
    env.execute("Autoscaling demo")

if __name__ == '__main__':
    logging.basicConfig(stream=sys.stdout, level=logging.INFO, format="%(message)s")
    autoscaling_demo()

```

To verify that rescale time estimation is working, make sure that DEBUG level logging of the Flink operator is enabled. The example below demonstrates how to update the helm chart file `values.yaml`. Then reinstall the updated helm chart and run your Flink job again.

```
log4j-operator.properties: |+
```

```
# Flink Operator Logging Overrides
rootLogger.level = DEBUG
```

Get the name of your leader pod.

```
ip=$(kubectl get configmap -n $NAMESPACE <job-name>-cluster-config-map -o json | jq
-r ".data[\"org.apache.flink.k8s.leader.restserver\"]" | awk -F: '{print $2}' | awk
-F '/' '{print $3}')

kubectl get pods -n $NAMESPACE -o json | jq -r ".items[] | select(.status.podIP ==
\"$ip\") | .metadata.name"
```

Run the following command to get the actual restart time used in metrics evaluations.

```
kubectl logs <FLINK-OPERATOR-POD-NAME> -c flink-kubernetes-operator -n <OPERATOR-
NAMESPACE> -f | grep "Restart time used in scaling summary computation"
```

You should see logs similar to the following. Note that only the first scaling uses `job.autoscaler.restart.time`. Subsequent scalings use the observed restart time.

```
2024-05-16 17:17:32,590 o.a.f.a.ScalingExecutor [DEBUG][default/autoscaler-
example] Restart time used in scaling summary computation: PT2M
2024-05-16 17:19:03,787 o.a.f.a.ScalingExecutor [DEBUG][default/autoscaler-
example] Restart time used in scaling summary computation: PT14S
2024-05-16 17:19:18,976 o.a.f.a.ScalingExecutor [DEBUG][default/autoscaler-
example] Restart time used in scaling summary computation: PT14S
2024-05-16 17:20:50,283 o.a.f.a.ScalingExecutor [DEBUG][default/autoscaler-
example] Restart time used in scaling summary computation: PT14S
2024-05-16 17:22:21,691 o.a.f.a.ScalingExecutor [DEBUG][default/autoscaler-
example] Restart time used in scaling summary computation: PT14S
```

7.0.0 and 7.1.0

The open source built-in Flink Autoscaler uses numerous metrics to make the best scaling decisions. However, the default values it uses for its calculations are meant to be applicable to most workloads and might not be optimal for a given job. The autotuning feature added into the Amazon EMR on EKS version of the Flink Operator looks at historical trends observed over specific captured metrics and then accordingly tries to calculate the most optimal value tailored for the given job.

Configuration	Required	Default	Description
kubernetes.operator.job.autoscaler.autotune.enable	No	False	Indicates whether the Flink Autoscaler should automatically tune configurations over time to optimize autoscalers scaling decisions. Currently , the Autoscaler can only autotune the Autoscaler parameter <code>restart.time</code> .
kubernetes.operator.job.autoscaler.autotune.metrics.history.max.count	No	3	Indicates how many historical Amazon EMR on EKS metrics the Autoscaler keeps in the Amazon EMR on EKS metrics config map.
kubernetes.operator.job.autoscaler.autotune.metrics.restart.count	No	3	Indicates how many number of restarts the Autoscaler performs before it starts calculating the average restart time for a given job.

To enable autotuning, you must have completed the following:

- Set `kubernetes.operator.job.autoscaler.autotune.enable`: to `true`
- Set `metrics.job.status.enable`: to `TOTAL_TIME`
- Followed the setup of [Using Autoscaler for Flink applications](#) to enable Autoscaling.

The following is an example deployment spec you can use to try out autotuning.

```
apiVersion: flink.apache.org/v1beta1
kind: FlinkDeployment
metadata:
  name: autoscaling-example
spec:
  flinkVersion: v1_18
```

flinkConfiguration:

```

# Autotuning parameters
kubernetes.operator.job.autoscaler.autotune.enable: "true"
kubernetes.operator.job.autoscaler.autotune.metrics.history.max.count: "2"
kubernetes.operator.job.autoscaler.autotune.metrics.restart.count: "1"
metrics.job.status.enable: TOTAL_TIME

# Autoscaler parameters
kubernetes.operator.job.autoscaler.enabled: "true"
kubernetes.operator.job.autoscaler.scaling.enabled: "true"
kubernetes.operator.job.autoscaler.stabilization.interval: "5s"
kubernetes.operator.job.autoscaler.metrics.window: "1m"

jobmanager.scheduler: adaptive

taskmanager.numberOfTaskSlots: "1"
state.savepoints.dir: s3://<S3_bucket>/autoscaling/savepoint/
state.checkpoints.dir: s3://<S3_bucket>/flink/autoscaling/checkpoint/
pipeline.max-parallelism: "4"

executionRoleArn: <JOB_ARN>
emrReleaseLabel: emr-6.14.0-flink-latest
jobManager:
  highAvailabilityEnabled: true
  storageDir: s3://<S3_bucket>/flink/autoscaling/ha/
  replicas: 1
  resource:
    memory: "1024m"
    cpu: 0.5
taskManager:
  resource:
    memory: "1024m"
    cpu: 0.5
job:
  jarURI: s3://<S3_bucket>/some-job-with-back-pressure
  parallelism: 1
  upgradeMode: last-state

```

To simulate backpressure, use the following deployment spec.

```

job:
  jarURI: s3://<S3_bucket>/pyflink-script.py

```

```
entryClass: "org.apache.flink.client.python.PythonDriver"
args: ["-py", "/opt/flink/usrlib/pyflink-script.py"]
parallelism: 1
upgradeMode: last-state
```

Upload the following Python script to your S3 bucket.

```
import logging
import sys
import time
import random

from pyflink.datastream import StreamExecutionEnvironment
from pyflink.table import StreamTableEnvironment

TABLE_NAME="orders"
QUERY=f"""
CREATE TABLE {TABLE_NAME} (
  id INT,
  order_time AS CURRENT_TIMESTAMP,
  WATERMARK FOR order_time AS order_time - INTERVAL '5' SECONDS
)
WITH (
  'connector' = 'datagen',
  'rows-per-second'='10',
  'fields.id.kind'='random',
  'fields.id.min'='1',
  'fields.id.max'='100'
);
"""

def create_backpressure(i):
    time.sleep(2)
    return i

def autoscaling_demo():
    env = StreamExecutionEnvironment.get_execution_environment()
    t_env = StreamTableEnvironment.create(env)
    t_env.execute_sql(QUERY)
    res_table = t_env.from_path(TABLE_NAME)

    stream = t_env.to_data_stream(res_table) \
        .shuffle().map(lambda x: create_backpressure(x))\
```

```

        .print()
        env.execute("Autoscaling demo")

if __name__ == '__main__':
    logging.basicConfig(stream=sys.stdout, level=logging.INFO, format="%(message)s")
    autoscaling_demo()

```

To verify that your autotuner is working, use the following commands. Note that you must use your own leader pod information for the Flink Operator.

First get the name of your leader pod.

```

ip=$(kubectl get configmap -n $NAMESPACE <job-name>-cluster-config-map -o json | jq
-r ".data[\"org.apache.flink.k8s.leader.restserver\"]" | awk -F: '{print $2}' | awk
-F '/' '{print $3}')

kubectl get pods -n $NAMESPACE -o json | jq -r ".items[] | select(.status.podIP ==
\"$ip\") | .metadata.name"

```

Once you have the name of your leader pod, you can run the following command.

```

kubectl logs -n $NAMESPACE -c flink-kubernetes-operator --follow <YOUR-FLINK-
OPERATOR-POD-NAME> | grep -E 'EmrEks|autotun|calculating|restart|autoscaler'

```

You should see logs similar to the following.

```

[m[33m2023-09-13 20:10:35,941[m [36mc.a.c.f.k.o.a.EmrEksMetricsAutotuner[m
[36m[DEBUG][flink/autoscaling-example] Using the latest
Emr Eks Metric for calculating restart.time for autotuning:
EmrEksMetrics(restartMetric=RestartMetric(restartingTime=65, numRestarts=1))

[m[33m2023-09-13 20:10:35,941[m [36mc.a.c.f.k.o.a.EmrEksMetricsAutotuner[m
[32m[INFO ][flink/autoscaling-example] Calculated average restart.time metric via
autotuning to be: PT0.065S

```

Maintenance and troubleshooting for Flink jobs on Amazon EMR on EKS

The following sections outline how to maintain your long-running Flink jobs, and provide guidance on how to troubleshoot some common issues with Flink jobs.

Maintaining Flink applications

Topics

- [Upgrade modes](#)

Flink applications are typically designed to run for long periods of time such as weeks, months, or even years. As with all long-running services, Flink streaming applications need to be maintained. This includes bug fixes, improvements, and migration to a Flink cluster of a later version.

When the spec changes for `FlinkDeployment` and `FlinkSessionJob` resources, you need to upgrade the running application. To do this, the operator stops the running job (unless already suspended) and redeploys it with the latest spec and, for stateful applications, the state from the previous run.

Users control how to manage the state when stateful applications stop and restore with the `upgradeMode` setting of the `JobSpec`.

Upgrade modes

Optional introduction

Stateless

Stateless application upgrades from empty state.

Last state

Quick upgrades in any application state (even for failing jobs), does not require a healthy job as it always uses the latest successful checkpoint. Manual recovery may be necessary if HA metadata is lost. To limit the time the job may fall back when picking up the latest checkpoint you can configure `kubernetes.operator.job.upgrade.last-state.max.allowed.checkpoint.age`. If the checkpoint is older than the configured value, a savepoint will be taken instead for healthy jobs. This is not supported in Session mode.

Savepoint

Use savepoint for upgrade, providing maximal safety and possibility to serve as backup/fork point. The savepoint will be created during the upgrade process. Note that the Flink job needs to be running to allow the savepoint to get created. If the job is in an unhealthy state, the last checkpoint will be used (unless `kubernetes.operator.job.upgrade.last-state-fallback.enabled` is set to false). If the last checkpoint is not available, the job upgrade will fail.

Troubleshooting

This section describes how to troubleshoot problems with Amazon EMR on EKS. For information on how to troubleshoot general problems with Amazon EMR, see [Troubleshoot a cluster](#) in the *Amazon EMR Management Guide*.

- [Troubleshooting jobs that use PersistentVolumeClaims \(PVC\)](#)
- [Troubleshooting Amazon EMR on EKS vertical autoscaling](#)
- [Troubleshooting Amazon EMR on EKS Spark operator](#)

Troubleshooting Apache Flink on Amazon EMR on EKS

Resource mapping not found when installing the Helm chart

You might encounter the following error message when you install the Helm chart.

```
Error: INSTALLATION FAILED: pulling from host 1234567890.dkr.ecr.us-west-2.amazonaws.com failed with status code [manifests 6.13.0]: 403 Forbidden Error:
INSTALLATION FAILED: unable to build kubernetes objects from release manifest:
[resource mapping not found for name: "flink-operator-serving-cert" namespace: "<the
namespace to install your operator>" from "": no matches for kind "Certificate" in
version "cert-manager.io/v1"

ensure CRDs are installed first, resource mapping not found for name: "flink-operator-
selfsigned-issuer" namespace: "<the namespace to install your operator>" " from "": no
matches for kind "Issuer" in version "cert-manager.io/v1"

ensure CRDs are installed first].
```

To resolve this error, install cert-manager to enable adding the webhook component. You must install cert-manager to each Amazon EKS cluster that you use.

```
kubectl apply -f https://github.com/cert-manager/cert-manager/releases/download/v1.12.0
```

AWS service access denied error

If you see an *access denied* error, confirm that the IAM role for `operatorExecutionRoleArn` in the Helm chart `values.yaml` file has the correct permissions. Also ensure the IAM role under `executionRoleArn` in your `FlinkDeployment` specification has the correct permissions.

FlinkDeployment is stuck

If your FlinkDeployment stalls in an arrested state, use the following steps to force delete the deployment:

1. Edit the deployment run.

```
kubectl edit -n Flink Namespace flinkdeployments/App Name
```

2. Remove this finalizer.

```
finalizers:
  - flinkdeployments.flink.apache.org/finalizer
```

3. Delete the deployment.

```
kubectl delete -n Flink Namespace flinkdeployments/App Name
```

s3a AWSBadRequestException issue when running a Flink application in an opt-in AWS Region

If you run a Flink application in an [opt-in AWS Region](#), you might see the following errors:

```
Caused by: org.apache.hadoop.fs.s3a.AWSBadRequestException: getFileStatus on
s3://flink.txt: com.amazonaws.services.s3.model.AmazonS3Exception: Bad Request
(Service: Amazon S3; Status Code: 400; Error Code: 400 Bad Request; Request ID:
ABCDEFGHIJKL; S3 Extended Request ID:
ABCDEFGHIJKLMNOP=; Proxy: null), S3 Extended Request ID: ABCDEFGHIJKLMNOP=:400 Bad
Request: Bad Request
(Service: Amazon S3; Status Code: 400; Error Code: 400 Bad Request; Request ID:
ABCDEFGHIJKL; S3 Extended Request ID: ABCDEFGHIJKLMNOP=; Proxy: null)
```

```
Caused by: org.apache.hadoop.fs.s3a.AWSBadRequestException: getS3Region on flink-
application: software.amazon.awssdk.services.s3.model.S3Exception: null
(Service: S3, Status Code: 400, Request ID: ABCDEFGHIJKLMNOP, Extended Request ID:
ABCDEFGHIJKLMNOPQRST==):null: null
(Service: S3, Status Code: 400, Request ID: ABCDEFGHIJKLMNOP, Extended Request ID:
AH142uDNaTUF0us/5IIVNvSakBcMjMCH7dd37ky0vE6jhABCDEFGHIJKLMNOPQRST==)
```

To fix these errors, use the following configuration in your FlinkDeployment definition file.

```
spec:
```

```

flinkConfiguration:
  taskmanager.numberOfTaskSlots: "2"
  fs.s3a.endpoint.region: OPT_IN_AWS_REGION_NAME

```

We also recommend that you use the SDKv2 credentials provider:

```

fs.s3a.aws.credentials.provider:
  software.amazon.awssdk.auth.credentials.WebIdentityTokenFileCredentialsProvider

```

If you want to use the SDKv1 credentials provider, make sure that your SDK supports your opt-in Region. For more information, see the [aws-sdk-java GitHub repository](#).

If you get `S3 AWSBadRequestException` when you run Flink SQL statements in an opt-in Region, make sure that you set the configuration `fs.s3a.endpoint.region`: *OPT_IN_AWS_REGION_NAME* in your flink configuration spec.

S3A AWSBadRequestException when running a Flink session job in CN regions

For Amazon EMR releases 6.15.0 - 7.2.0, you might encounter the following error messages when you run a Flink session job in CN regions. These include China (Beijing) and China (Ningxia):

```

Error:
{"type":"org.apache.flink.kubernetes.operator.exception.ReconciliationException","message":"or
      getFileStatus on s3://ABCDPath:
software.amazon.awssdk.services.s3.model.S3Exception: null (Service: S3, Status Code:
400, Request ID: ABCDEFGH, Extended Request ID:
      ABCDEFGH:null: null (Service: S3, Status Code: 400, Request ID:
ABCDEFGH, Extended Request ID: ABCDEFGH","additionalMetadata":{},"throwableList":

[{"type":"org.apache.hadoop.fs.s3a.AWSBadRequestException","message":"getFileStatus on
s3://ABCDPath: software.amazon.awssdk.services.s3.model.S3Exception:
      null (Service: S3, Status Code: 400, Request ID: ABCDEFGH, Extended
Request ID: ABCDEFGH:null: null (Service: S3, Status Code: 400, Request ID: ABCDEFGH,
      Extended Request ID: ABCDEFGH","additionalMetadata":{}},
{"type":"software.amazon.awssdk.services.s3.model.S3Exception","message":"null
(Service: S3, Status Code: 400,
      Request ID: ABCDEFGH, Extended Request ID:
ABCDEFGH","additionalMetadata":{}}}]

```

There is an awareness of this issue. The team is working on patching the flink operators for all of these release versions. However, before we finish the patch, to fix this error, you need to download

the flink operator helm chart, untar it (extract the compressed file) and make configuration changes in the helm chart.

The specific steps are the following:

1. Change to, specifically change directories to, your local folder for the helm chart, and run the following command line to pull the helm chart and untar (extract) it.

```
helm pull oci://public.ecr.aws/emr-on-eks/flink-kubernetes-operator \
--version $VERSION \
--namespace $NAMESPACE
```

```
tar -zxvf flink-kubernetes-operator-$VERSION.tgz
```

2. Go into the helm chart folder and find the templates/flink-operator.yaml file.
3. Find the flink-operator-config ConfigMap and add the following fs.s3a.endpoint.region configuration in the flink-conf.yaml. For example:

```
{{- if .Values.defaultConfiguration.create }}
apiVersion: v1
kind: ConfigMap
metadata:
  name: flink-operator-config
  namespace: {{ .Release.Namespace }}
  labels:
    {{- include "flink-operator.labels" . | nindent 4 }}
data:
  flink-conf.yaml: |+
fs.s3a.endpoint.region: {{ .Values.emrContainers.awsRegion }}
```

4. Install the local helm chart and run your job.

Supported releases for Amazon EMR on EKS with Apache Flink

Apache Flink is available with the following Amazon EMR on EKS releases. For information on all of the releases that are available, see [Amazon EMR on EKS releases](#).

Release label	Java	Flink	Flink operator
emr-7.2.0-flink-latest	17	1.18.1	-
emr-7.2.0-flink-k8s-operator-latest	11	-	1.8.0
emr-7.1.0-flink-latest	17	1.18.1	-
emr-7.1.0-flink-k8s-operator-latest	11	-	1.6.1
emr-7.0.0-flink-latest	11	1.18.0	-
emr-7.0.0-flink-k8s-operator-latest	11	-	1.6.1
emr-6.15.0-flink-latest	11	1.17.1	-
emr-6.15.0-flink-k8s-operator-latest	11	-	1.6.0
emr-6.14.0-flink-latest	11	1.17.1	-
emr-6.14.0-flink-k8s-operator-latest	11	-	1.6.0
emr-6.13.0-flink-latest	11	1.17.0	-
emr-6.13.0-flink-k8s-operator-latest	11	-	1.5.0

Running Spark jobs with Amazon EMR on EKS

A *job run* is a unit of work, such as a Spark jar, PySpark script, or SparkSQL query, that you submit to Amazon EMR on EKS. This topic provides an overview of managing job runs using the AWS CLI, viewing job runs using the Amazon EMR console, and troubleshooting common job run errors.

Note that you can't run IPv6 Spark jobs on Amazon EMR on EKS

Note

Before you submit a job run with Amazon EMR on EKS, you must complete the steps in [Setting up Amazon EMR on EKS](#).

Topics

- [Running Spark jobs with StartJobRun](#)
- [Running Spark jobs with the Spark operator](#)
- [Running Spark jobs with spark-submit](#)
- [Using Apache Livy with Amazon EMR on EKS](#)
- [Managing Amazon EMR on EKS job runs](#)
- [Using job templates](#)
- [Using pod templates](#)
- [Using job retry policies](#)
- [Using Spark event log rotation](#)
- [Using Spark container log rotation](#)
- [Using vertical autoscaling with Amazon EMR Spark jobs](#)

Running Spark jobs with StartJobRun

This section includes detailed setup steps to get your environment ready to run Spark jobs and then provides step-by-step instructions for submitting a job run with specified parameters.

Topics

- [Setting up Amazon EMR on EKS](#)

- [Submit a job run with StartJobRun](#)
- [Using job submitter classification](#)
- [Using Amazon EMR container defaults classification](#)

Setting up Amazon EMR on EKS

Complete the following tasks to get set up for Amazon EMR on EKS. If you've already signed up for Amazon Web Services (AWS) and have been using Amazon EKS, you are almost ready to use Amazon EMR on EKS. Skip any of the tasks that you've already completed.

Note

You can also follow the [Amazon EMR on EKS Workshop](#) to set up all the necessary resources to run Spark jobs on Amazon EMR on EKS. The workshop also provides automation by using CloudFormation templates to create the resources necessary for you to get started. For other templates and best practices, see our [EMR Containers Best Practices Guide](#) on GitHub.

1. [Install or update to the latest version of the AWS CLI](#)
2. [Set up kubectl and eksctl](#)
3. [Get started with Amazon EKS – eksctl](#)
4. [Enable cluster access for Amazon EMR on EKS](#)
5. [Enable IAM Roles for the EKS cluster](#)
6. [Grant users access to Amazon EMR on EKS](#)
7. [Register the Amazon EKS cluster with Amazon EMR](#)

Enable cluster access for Amazon EMR on EKS

The following sections show a couple ways to enable cluster access. The first is by using Amazon EKS cluster access management (CAM) and the latter shows how to take manual steps to enable cluster access.

Enable cluster access using EKS Access Entry (recommended)

Note

The `aws-auth` ConfigMap is deprecated. The recommended method to manage access to Kubernetes APIs is [Access Entries](#).

Amazon EMR is integrated with [Amazon EKS cluster access management \(CAM\)](#), so you can automate configuration of the necessary AuthN and AuthZ policies to run Amazon EMR Spark jobs in namespaces of Amazon EKS clusters. When you create a virtual cluster from an Amazon EKS cluster namespace, Amazon EMR automatically configures all of the necessary permissions, so you don't need to add any extra steps into your current workflows.

Note

The Amazon EMR integration with Amazon EKS CAM is supported only for new Amazon EMR on EKS virtual clusters. You can't migrate existing virtual clusters to use this integration.

Prerequisites

- Make sure that you are running version 2.15.3 or higher of the AWS CLI
- Your Amazon EKS cluster must be on version 1.23 or higher.

Setup

To set up the integration between Amazon EMR and the AccessEntry API operations from Amazon EKS, make sure that you have completed the follow items:

- Make sure that `authenticationMode` of your Amazon EKS cluster is set to `API_AND_CONFIG_MAP`.

```
aws eks describe-cluster --name <eks-cluster-name>
```

If it isn't already, set `authenticationMode` to `API_AND_CONFIG_MAP`.

```
aws eks update-cluster-config
  --name <eks-cluster-name>
  --access-config authenticationMode=API_AND_CONFIG_MAP
```

For more information about authentication modes, see [Cluster authentication modes](#).

- Make sure that the [IAM role](#) that you're using to run the `CreateVirtualCluster` and `DeleteVirtualCluster` API operations also has the following permissions:

```
{
  "Effect": "Allow",
  "Action": [
    "eks:CreateAccessEntry"
  ],
  "Resource":
    "arn:<AWS_PARTITION>:eks:<AWS_REGION>:<AWS_ACCOUNT_ID>:cluster/<EKS_CLUSTER_NAME>"
},
{
  "Effect": "Allow",
  "Action": [
    "eks:DescribeAccessEntry",
    "eks>DeleteAccessEntry",
    "eks>ListAssociatedAccessPolicies",
    "eks:AssociateAccessPolicy",
    "eks:DisassociateAccessPolicy"
  ],
  "Resource": "arn:<AWS_PARTITION>:eks:<AWS_REGION>:<AWS_ACCOUNT_ID>:access-entry/
    <EKS_CLUSTER_NAME>/role/<AWS_ACCOUNT_ID>/AWSServiceRoleForAmazonEMRContainers/*"
}
```

- To create a least permissive policy, add the context keys `accessScope` and `namespaces`. The following example shows a policy with these context keys:

```
{
  "Effect": "Allow",
  "Action": [
    "eks:CreateAccessEntry"
  ],
  "Resource": [
    "arn:<AWS_PARTITION>:eks:<AWS_REGION>:<AWS_ACCOUNT_ID>:cluster/<EKS_CLUSTER_NAME>"
  ]
}
```

```

    ],
    {
      "Effect": "Allow",
      "Action": [
        "eks:DescribeAccessEntry",
        "eks:DeleteAccessEntry",
        "eks:ListAssociatedAccessPolicies"
      ],
      "Resource": [
        "arn:<AWS_PARTITION>:eks:<AWS_REGION>:<AWS_ACCOUNT_ID>:access-
entry/<EKS_CLUSTER_NAME>/role/<AWS_ACCOUNT_ID>/AWSServiceRoleForAmazonEMRContainers/
*"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "eks:DisassociateAccessPolicy",
        "eks:AssociateAccessPolicy"
      ],
      "Resource": [
        "arn:<AWS_PARTITION>:eks:<AWS_REGION>:<AWS_ACCOUNT_ID>:access-
entry/<EKS_CLUSTER_NAME>/role/<AWS_ACCOUNT_ID>/AWSServiceRoleForAmazonEMRContainers/
*"
      ],
      "Condition": {
        "StringEquals": {
          "eks:accessScope": "namespace"
        },
        "ForAllValues:StringEquals": {
          "eks:namespaces": [
            "<EKS_NAMESPACE>"
          ]
        }
      }
    }
  ]
}

```

Concepts and terminology

The following is a list of terminologies and concepts related to Amazon EKS CAM.

- Virtual cluster (VC) – logical representation of the namespace created in Amazon EKS. It's a 1:1 link to an Amazon EKS cluster namespace. You can use it to run Amazon EMR workloads on a Amazon EKS cluster within the specified namespace.
- Namespace – mechanism to isolate groups of resources within a single EKS cluster.
- Access policy – permissions that grant access and actions to an IAM role within an EKS cluster.
- Access entry – an entry created with a role arn. You can link the access entry to an access policy to assign specific permissions in the Amazon EKS cluster.
- EKS access entry integrated virtual cluster – the virtual cluster created using [access entry API operations](#) from Amazon EKS.

Enable cluster access using aws-auth

You must allow Amazon EMR on EKS access to a specific namespace in your cluster by taking the following actions: creating a Kubernetes role, binding the role to a Kubernetes user, and mapping the Kubernetes user with the service linked role [AWSServiceRoleForAmazonEMRContainers](#). These actions are automated in `eksctl` when the IAM identity mapping command is used with `emr-containers` as the service name. You can perform these operations easily by using the following command.

```
eksctl create iamidentitymapping \  
  --cluster my_eks_cluster \  
  --namespace kubernetes_namespace \  
  --service-name "emr-containers"
```

Replace *my_eks_cluster* with the name of your Amazon EKS cluster and replace *kubernetes_namespace* with the Kubernetes namespace created to run Amazon EMR workloads.

Important

You must download the latest `eksctl` using the previous step [Set up kubectl and eksctl](#) to use this functionality.

Manual steps to enable cluster access for Amazon EMR on EKS

You can also use the following manual steps to enable cluster access for Amazon EMR on EKS.

1. Create a Kubernetes role in a specific namespace

Amazon EKS 1.22 - 1.29

With Amazon EKS 1.22 - 1.29, run the following command to create a Kubernetes role in a specific namespace. This role grants the necessary RBAC permissions to Amazon EMR on EKS.

```
namespace=my-namespace
cat - >>EOF | kubectl apply -f - >>namespace "${namespace}"
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: emr-containers
  namespace: ${namespace}
rules:
  - apiGroups: [""]
    resources: ["namespaces"]
    verbs: ["get"]
  - apiGroups: [""]
    resources: ["serviceaccounts", "services", "configmaps", "events", "pods",
"pods/log"]
    verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"deletcollection", "annotate", "patch", "label"]
  - apiGroups: [""]
    resources: ["secrets"]
    verbs: ["create", "patch", "delete", "watch"]
  - apiGroups: ["apps"]
    resources: ["statefulsets", "deployments"]
    verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"annotate", "patch", "label"]
  - apiGroups: ["batch"]
    resources: ["jobs"]
    verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"annotate", "patch", "label"]
  - apiGroups: ["extensions", "networking.k8s.io"]
    resources: ["ingresses"]
    verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"annotate", "patch", "label"]
  - apiGroups: ["rbac.authorization.k8s.io"]
    resources: ["roles", "rolebindings"]
    verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"deletcollection", "annotate", "patch", "label"]
  - apiGroups: [""]
```

```
resources: ["persistentvolumeclaims"]
verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"deletecollection", "annotate", "patch", "label"]
EOF
```

Amazon EKS 1.21 and below

With Amazon EKS 1.21 and below, run the following command to create a Kubernetes role in a specific namespace. This role grants the necessary RBAC permissions to Amazon EMR on EKS.

```
namespace=my-namespace
cat - >>EOF | kubectl apply -f - >>namespace "${namespace}"
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: emr-containers
  namespace: ${namespace}
rules:
- apiGroups: [""]
  resources: ["namespaces"]
  verbs: ["get"]
- apiGroups: [""]
  resources: ["serviceaccounts", "services", "configmaps", "events", "pods",
"pods/log"]
  verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"deletecollection", "annotate", "patch", "label"]
- apiGroups: [""]
  resources: ["secrets"]
  verbs: ["create", "patch", "delete", "watch"]
- apiGroups: ["apps"]
  resources: ["statefulsets", "deployments"]
  verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"annotate", "patch", "label"]
- apiGroups: ["batch"]
  resources: ["jobs"]
  verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"annotate", "patch", "label"]
- apiGroups: ["extensions"]
  resources: ["ingresses"]
  verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"annotate", "patch", "label"]
```

```

- apiGroups: ["rbac.authorization.k8s.io"]
  resources: ["roles", "rolebindings"]
  verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"deletecollection", "annotate", "patch", "label"]
- apiGroups: [""]
  resources: ["persistentvolumeclaims"]
  verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"deletecollection", "annotate", "patch", "label"]
EOF

```

2. Create a Kubernetes role binding scoped to the namespace

Run the following command to create a Kubernetes role binding in the given namespace. This role binding grants the permissions defined in the role created in the previous step to a user named `emr-containers`. This user identifies [service-linked roles for Amazon EMR on EKS](#) and thus allows Amazon EMR on EKS to perform actions as defined by the role you created.

```

namespace=my-namespace

cat - <<EOF | kubectl apply -f - --namespace "${namespace}"
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: emr-containers
  namespace: ${namespace}
subjects:
- kind: User
  name: emr-containers
  apiGroup: rbac.authorization.k8s.io
roleRef:
  kind: Role
  name: emr-containers
  apiGroup: rbac.authorization.k8s.io
EOF

```

3. Update Kubernetes `aws-auth` configuration map

You can use one of the following options to map the Amazon EMR on EKS service-linked role with the `emr-containers` user that was bound with the Kubernetes role in the previous step.

Option 1: Using `eksctl`

Run the following `eksctl` command to map the Amazon EMR on EKS service-linked role with the `emr-containers` user.

```
eksctl create iamidentitymapping \  
  --cluster my-cluster-name \  
  --arn "arn:aws:iam::my-account-id:role/AWSServiceRoleForAmazonEMRContainers" \  
  --username emr-containers
```

Option 2: Without using eksctl

1. Run the following command to open the `aws-auth` configuration map in text editor.

```
kubectl edit -n kube-system configmap/aws-auth
```

Note

If you receive an error stating `Error from server (NotFound): configmaps "aws-auth" not found`, see the steps in [Add user roles](#) in the Amazon EKS User Guide to apply the stock ConfigMap.

2. Add Amazon EMR on EKS service-linked role details to the `mapRoles` section of the ConfigMap, under `data`. Add this section if it does not already exist in the file. The updated `mapRoles` section under `data` looks like the following example.

```
apiVersion: v1  
data:  
  mapRoles: |  
    - rolearn: arn:aws:iam::<your-account-id>:role/  
    AWSServiceRoleForAmazonEMRContainers  
      username: emr-containers  
    - ... <other previously existing role entries, if there's any>.
```

3. Save the file and exit your text editor.

Enable cluster access for Amazon SageMaker Unified Studio

Amazon EMR on EKS and Amazon SageMaker Unified Studio require access to an Amazon EKS cluster. Please follow the steps at [Enable EKS cluster access for EMR on EKS and SageMaker Unified Studio](#) to provide access.

Enable IAM Roles for the EKS cluster

The following topics detail options for enabling IAM roles.

Topics

- [Option 1: Enable EKS Pod Identity on the EKS Cluster](#)
- [Option 2: Enable IAM Roles for Service Accounts \(IRSA\) on the EKS cluster](#)

Option 1: Enable EKS Pod Identity on the EKS Cluster

Amazon EKS Pod Identity associations provide the ability to manage credentials for your applications, similar to the way that Amazon EC2 instance profiles provide credentials to Amazon EC2 instances. Amazon EKS Pod Identity provides credentials to your workloads with an additional EKS Auth API and an agent pod that runs on each node.

Amazon EMR on EKS starts to support EKS pod identity since emr-7.3.0 release for the StartJobRun submission model.

For more information on EKS pod identities, refer to [Understand how EKS Pod Identity works](#).

Why EKS Pod Identities?

As part of EMR setup, the Job Execution Role needs to establish trust boundaries between an IAM role and service accounts in a specific namespace (of EMR virtual clusters). With IRSA, this was achieved by updating the trust policy of the EMR Job Execution Role. However, due to the 4096 character hard-limit on IAM trust policy length, there was a constraint to share a single Job Execution IAM Role across a maximum of twelve (12) EKS clusters.

With EMR's support for Pod Identities, the trust boundary between IAM roles and service accounts are now being managed by the EKS team through EKS pod identity's association APIs.

Note

The security boundary for EKS pod identity is still on service account level, not on pod level.

Pod Identity Considerations

For information on the Pod Identity Limitations, see [EKS Pod Identity considerations](#).

Prepare EKS Pod Identity in EKS Cluster

Check if the required permission exists in `NodeInstanceRole`

The node role `NodeInstanceRole` needs a permission for the agent to do the `AssumeRoleForPodIdentity` action in the EKS Auth API. You can add the following to the [AmazonEKSWorkerNodePolicy](#), which is defined in the Amazon EKS User Guide, or use a custom policy.

If your EKS cluster was created with `eksctl` version higher than **0.181.0**, the `AmazonEKSWorkerNodePolicy`, including the required `AssumeRoleForPodIdentity` permission, will be attached to the node role automatically. If the permission is not present, manually add the following permission to `AmazonEKSWorkerNodePolicy` that allows assuming a role for pod identity. This permission is needed by the EKS pod identity agent to retrieve credentials for pods.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "eks-auth:AssumeRoleForPodIdentity"
      ],
      "Resource": [
        "*"
      ],
      "Sid": "AllowEKSAUTHAssumeroleforpodidentity"
    }
  ]
}
```

```
]
}
```

Create EKS pod identity agent add-on

Use the following command to create EKS Pod Identity Agent add-on with the latest version:

```
aws eks create-addon --cluster-name cluster-name --addon-name eks-pod-identity-agent
kubect1 get pods -n kube-system | grep 'eks-pod-identity-agent'
```

Use the following steps to create EKS Pod Identity Agent add-on from the Amazon EKS console:

1. Open the Amazon EKS console: [Amazon EKS console](#).
2. In the left navigation pane, select **Clusters**, and then select the name of the cluster that you want to configure the EKS Pod Identity Agent add-on for.
3. Choose the **Add-ons** tab.
4. Choose **Get more add-ons**.
5. Select the box in the top right of the add-on box for EKS Pod Identity Agent and then choose **Next**.
6. On the **Configure selected add-ons settings** page, select any version in the **Version** drop-down list.
7. (Optional) Expand **Optional configuration settings** to enter additional configuration. For example, you can provide an alternative container image location and ImagePullSecrets. The JSON Schema with accepted keys is shown in **Add-on configuration schema**.

Enter the configuration keys and values in **Configuration values**.

8. Choose **Next**.
9. Confirm that the agent pods are running on your cluster via the CLI.

```
kubect1 get pods -n kube-system | grep 'eks-pod-identity-agent'
```

An example output is as followings:

NAME	READY	STATUS	RESTARTS	AGE
eks-pod-identity-agent-gmqp7	1/1	Running	1 (24h ago)	24h

eks-pod-identity-agent-prnsh	1/1	Running	1 (24h ago)	24h
------------------------------	-----	---------	-------------	-----

This sets up a new DaemonSet in the kube-system namespace. The Amazon EKS Pod Identity Agent, running on each EKS node, uses the [AssumeRoleForPodIdentity](#) action to retrieve temporary credentials from the EKS Auth API. These credentials are then made available for the AWS SDKs that you run inside your containers.

For more information, check the pre-requisite in the public document: [Set up the Amazon EKS Pod Identity Agent](#).

Create a Job Execution Role

Create or update job execution role that allows EKS Pod Identity

To run workloads with Amazon EMR on EKS, you need to create an IAM role. We refer to this role as the job execution role in this documentation. For more information about how to create the IAM role, see [Creating IAM roles](#) in the user Guide.

Additionally, you must create an IAM policy that specifies the necessary permissions for the job execution role and then attach this policy to the role to enable EKS Pod Identity.

For example, you have the following job execution role. For more information, see [Create a job execution role](#).

```
arn:aws:iam::111122223333:role/PodIdentityJobExecutionRole
```

Important

Amazon EMR on EKS automatically creates Kubernetes Service Accounts, based on your job execution role name. Ensure the role name is not too long, as your job may fail if the combination of `cluster_name`, `pod_name`, and `service_account_name` exceeds the length limit.

Job Execution Role Configuration – Ensure the job execution role is created with the below trust permission for EKS Pod Identity. To update an existing job execution role, configure it to trust the following EKS service principal as an additional permission in the trust policy. This trust permission can co-exist with existing IRSA trust policies.


```
cat >trust-relationship.json <<EOF
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowEksAuthToAssumeRoleForPodIdentity",
      "Effect": "Allow",
      "Principal": {
        "Service": "pods.eks.amazonaws.com"
      },
      "Action": [
        "sts:AssumeRole",
        "sts:TagSession"
      ]
    }
  ]
}
EOF
```

User Permission: Users require the `iam:PassRole` permission to execute `StartJobRun` API calls or submit jobs. This permission enables users to pass the job execution role to EMR on EKS. Job administrators should have the permission by default.

Below is the permission needed for a user:

```
{
  "Effect": "Allow",
  "Action": "iam:PassRole",
  "Resource": "arn:aws:iam::111122223333:role/PodIdentityJobExecutionRole",
  "Condition": {
    "StringEquals": {
      "iam:PassedToService": "pods.eks.amazonaws.com"
    }
  }
}
```

To further restrict the user access to specific EKS clusters, add the `AssociatedResourceArn` attribute filter to the IAM policy. It limits the role assumption to authorized EKS clusters, strengthening your resource-level security controls.

```
"Condition": {
  "ArnLike": {
```

```
    "iam:AssociatedResourceARN": [
      "arn:aws:eks:us-west-2:111122223333:cluster/*"
    ]
  }
}
```

Set up EKS pod identity associations

Prerequisite

Make sure the IAM Identity creating the pod identity association, such as an EKS admin user, has the permission `eks:CreatePodIdentityAssociation` and `iam:PassRole`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "eks:CreatePodIdentityAssociation"
      ],
      "Resource": [
        "arn:aws:eks:*:*:cluster/*"
      ],
      "Sid": "AllowEKSCreatepodidentityassociation"
    },
    {
      "Effect": "Allow",
      "Action": [
        "iam:PassRole"
      ],
      "Resource": [
        "arn:aws:iam:*:*:role/*"
      ],
      "Condition": {
        "StringEquals": {
          "iam:PassedToService": "pods.eks.amazonaws.com"
        }
      },
      "Sid": "AllowIAMPassrole"
    }
  ]
}
```

```
}
```

Create Associations for the role and EMR service account

Create EMR role associations through the AWS CLI

When you submit a job to a Kubernetes namespace, an administrator must create associations between the job execution role and the identity of the EMR managed service account. Note that the EMR managed service account is automatically created at job submission, scoped to the namespace where the job is submitted.

With the AWS CLI (above version 2.24.0), run the following command to create role associations with pod identity.

Run the following command to create role associations with pod identity:

```
aws emr-containers create-role-associations \
    --cluster-name mycluster \
    --namespace mynamespace \
    --role-name JobExecutionRoleIRSAv2
```

Note:

- Each cluster can have a limit of 1,000 associations. Each job execution role - namespace mapping will require 3 associations for job submitter, driver and executor pods.
- You can only associate roles that are in the same AWS account as the cluster. You can delegate access from another account to the role in this account that you configure for EKS Pod Identities to use. For a tutorial about delegating access and AssumeRole, see [IAM tutorial: Delegate access across AWS accounts using IAM roles](#).

Create EMR role associations through Amazon EKS

EMR creates service account with certain naming pattern when a job is submitted. To make manual associations or integrate this workflow with the AWS SDK, follow these steps:

Construct Service Account Name:

```
emr-containers-sa-spark-%(SPARK_ROLE)s-%(AWS_ACCOUNT_ID)s-
%(BASE36_ENCODED_ROLE_NAME)s
```

The below examples creates a role associations for a sample Job execution role JobExecutionRoleIRSAv2.

Example Role Associations:

```
RoleName: JobExecutionRoleIRSAv2
Base36EncodingOfRoleName: 2eum5fah1jc1kwyjc19ikdhdkdegh1n26vbe
```

Sample CLI command:

```
# setup for the client service account (used by job runner pod)
# emr-containers-sa-spark-client-111122223333-2eum5fah1jc1kwyjc19ikdhdkdegh1n26vbe
aws eks create-pod-identity-association --cluster-name mycluster
--role-arn arn:aws:iam::111122223333:role/JobExecutionRoleIRSAv2
--namespace mynamespace --service-account emr-containers-sa-spark-
client-111122223333-2eum5fah1jc1kwyjc19ikdhdkdegh1n26vbe

# driver service account
# emr-containers-sa-spark-driver-111122223333-2eum5fah1jc1kwyjc19ikdhdkdegh1n26vbe

aws eks create-pod-identity-association --cluster-name mycluster
--role-arn arn:aws:iam::111122223333:role/JobExecutionRoleIRSAv2
--namespace mynamespace --service-account emr-containers-sa-spark-
driver-111122223333-2eum5fah1jc1kwyjc19ikdhdkdegh1n26vbe

# executor service account
# emr-containers-sa-spark-executor-111122223333-2eum5fah1jc1kwyjc19ikdhdkdegh1n26vbe
aws eks create-pod-identity-association --cluster-name mycluster
--role-arn arn:aws:iam::111122223333:role/JobExecutionRoleIRSAv2
--namespace mynamespace --service-account emr-containers-sa-spark-
executor-111122223333-2eum5fah1jc1kwyjc19ikdhdkdegh1n26vbe
```

Once you completed all the steps required for EKS pod identity, you can skip the following steps for IRSA setup:

- [Enable IAM Roles for Service Accounts \(IRSA\) on the EKS cluster](#)
- [Create a job execution role](#)
- [Update the trust policy of the job execution role](#)

You can skip directly to the following step: [Grant users access to Amazon EMR on EKS](#)

Delete Role Associations

Whenever you delete a virtual cluster or a job execution role and you no longer want to give access to EMR to its service accounts, you should delete the associations for the role. This is because EKS allows associations with non-existent resources (namespace and service account). Amazon EMR on EKS recommends deleting the associations if the namespace is deleted or the role is no longer in use, to free up space for other associations.

Note

The lingering associations could potentially impact your ability to scale if you don't delete them, as EKS has limitations on the number of associations you can create (soft limit: 1000 associations per cluster). You can list pod identity associations in a given namespace to check if you have any lingering associations that needs to be cleaned up:

```
aws eks list-pod-identity-associations --cluster-name mycluster --namespace mynamespace
```

With the AWS CLI (version 2.24.0 or higher), run the following `emr-containers` command to delete EMR's role associations:

```
aws emr-containers delete-role-associations \
    --cluster-name mycluster \
    --namespace mynamespace \
    --role-name JobExecutionRoleIRSAv2
```

Automatically Migrate Existing IRSA to Pod Identity

You can use the tool `eksctl` to migrate existing IAM Roles for Service Accounts (IRSA) to pod identity associations:

```
eksctl utils migrate-to-pod-identity \
    --cluster mycluster \
    --remove-oidc-provider-trust-relationship \
    --approve
```

Running the command without the `--approve` flag will only output a plan reflecting the migration steps, and no actual migration will occur.

Troubleshooting

My job failed with `NoClassDefinitionFound` or `ClassNotFound` Exception for Credentials Provider, or failed to get credentials provider.

EKS Pod Identity uses the Container Credentials Provider to retrieve the necessary credentials. If you have specified a custom credentials provider, ensure it is working correctly. Alternatively, make sure you are using a correct AWS SDK version that supports the EKS Pod Identity. For more information, refer to [Get started with Amazon EKS](#).

Job failed with the "Failed to Retrieve Credentials Due to [x] Size Limit" error shown in the `eks-pod-identity-agent` log.

EMR on EKS creates Kubernetes Service Accounts based on the job execution role name. If the role name is too long, EKS Auth will fail to retrieve credentials because the combination of `cluster_name`, `pod_name`, and `service_account_name` exceeds the length limit. Identify which component is taking up the most space and adjust the size accordingly.

Job failed with "Failed to Retrieve Credentials xxx" error shown in the `eks-pod-identity` log.

One possible cause of this issue could be that the EKS cluster is configured under private subnets without correctly configuring PrivateLink for the cluster. Check if your cluster is in a private network and configure AWS PrivateLink to address the issue. For detailed instructions, refer to [Get started with Amazon EKS](#).

Option 2: Enable IAM Roles for Service Accounts (IRSA) on the EKS cluster

The IAM roles for service accounts feature is available on Amazon EKS versions 1.14 and later and for EKS clusters that are updated to versions 1.13 or later on or after September 3rd, 2019. To use this feature, you can update existing EKS clusters to version 1.14 or later. For more information, see [Updating an Amazon EKS cluster Kubernetes version](#).

If your cluster supports IAM roles for service accounts, it has an [OpenID Connect](#) issuer URL associated with it. You can view this URL in the Amazon EKS console, or you can use the following AWS CLI command to retrieve it.

Important

You must use the latest version of the AWS CLI to receive the proper output from this command.

```
aws eks describe-cluster --name cluster_name --query "cluster.identity.oidc.issuer" --output text
```

The expected output is as follows.

```
https://oidc.eks.<region-code>.amazonaws.com/id/EXAMPLED539D4633E53DE1B716D3041E
```

To use IAM roles for service accounts in your cluster, you must create an OIDC identity provider using either [eksctl](#) or the [AWS Management Console](#).

To create an IAM OIDC identity provider for your cluster with eksctl

Check your eksctl version with the following command. This procedure assumes that you have installed eksctl and that your eksctl version is 0.32.0 or later.

```
eksctl version
```

For more information about installing or upgrading eksctl, see [Installing or upgrading eksctl](#).

Create your OIDC identity provider for your cluster with the following command. Replace *cluster_name* with your own value.

```
eksctl utils associate-iam-oidc-provider --cluster cluster_name --approve
```

To create an IAM OIDC identity provider for your cluster with the AWS Management Console

Retrieve the OIDC issuer URL from the Amazon EKS console description of your cluster, or use the following AWS CLI command.

Use the following command to retrieve the OIDC issuer URL from the AWS CLI.

```
aws eks describe-cluster --name <cluster_name> --query "cluster.identity.oidc.issuer" --output text
```

Use the following steps to retrieve the OIDC issuer URL from the Amazon EKS console.

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation panel, choose **Identity Providers**, and then choose **Create Provider**.

1. For **Provider Type**, choose **Choose a provider type**, and then choose **OpenID Connect**.
2. For **Provider URL**, paste the OIDC issuer URL for your cluster.
3. For Audience, type sts.amazonaws.com and choose **Next Step**.
3. Verify that the provider information is correct, and then choose **Create** to create your identity provider.

Create a job execution role

To run workloads on Amazon EMR on EKS, you need to create an IAM role. We refer to this role as the *job execution role* in this documentation. For more information about how to create IAM roles, see [Creating IAM roles](#) in the IAM user Guide.

You must also create an IAM policy that specifies the permissions for the job execution role and then attach the IAM policy to the job execution role.

The following policy for the job execution role allows access to resource targets, Amazon S3, and CloudWatch. These permissions are necessary to monitor jobs and access logs. To follow the same process using the AWS CLI:

Create IAM Role for job execution: Let's create the role that EMR will use for job execution. This is the role, EMR jobs will assume when they run on EKS.

```
cat <<EOF > ~/environment/emr-trust-policy.json
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "elasticmapreduce.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
EOF

aws iam create-role --role-name EMRContainers-JobExecutionRole --assume-role-policy-
document file://~/environment/emr-trust-policy.json
```


Next, we need to attach the required IAM policies to the role so it can write logs to s3 and cloudwatch.

```
cat <<EoF > ~/environment/EMRContainers-JobExecutionRole.json
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:ListBucket"
      ],
      "Resource": "arn:aws:s3:::amzn-s3-demo-bucket"
    },
    {
      "Effect": "Allow",
      "Action": [
        "logs:PutLogEvents",
        "logs:CreateLogStream",
        "logs:DescribeLogGroups",
        "logs:DescribeLogStreams"
      ],
      "Resource": [
        "arn:aws:logs:*:*:*"
      ]
    }
  ]
}
EoF
aws iam put-role-policy --role-name EMRContainers-JobExecutionRole --policy-name
EMR-Containers-Job-Execution --policy-document file://~/environment/EMRContainers-
JobExecutionRole.json
```

Note

Access should be appropriately scoped, not granted to all S3 objects in the job execution role.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::amzn-s3-demo-bucket"
      ],
      "Sid": "AllowS3Putobject"
    },
    {
      "Effect": "Allow",
      "Action": [
        "logs:PutLogEvents",
        "logs:CreateLogStream",
        "logs:DescribeLogGroups",
        "logs:DescribeLogStreams"
      ],
      "Resource": [
        "arn:aws:logs:*:*:*"
      ],
      "Sid": "AllowLOGSPutlogevents"
    }
  ]
}
```

For more information, see [Using job execution roles](#), [Configure a job run to use S3 logs](#), and [Configure a job run to use CloudWatch Logs](#).

Update the trust policy of the job execution role

When you use IAM Roles for Service Accounts (IRSA) to run jobs on a Kubernetes namespace, an administrator must create a trust relationship between the job execution role and the identity of the EMR managed service account. The trust relationship can be created by updating the trust policy of the job execution role. Note that the EMR managed service account is automatically created at job submission, scoped to the namespace where the job is submitted.

Run the following command to update the trust policy.

```
aws emr-containers update-role-trust-policy \
  --cluster-name cluster \
  --namespace namespace \
  --role-name iam_role_name_for_job_execution
```

For more information, see [Using job execution roles with Amazon EMR on EKS](#).

Important

The operator running the above command must have these permissions:
eks:DescribeCluster, iam:GetRole, iam:UpdateAssumeRolePolicy.

Grant users access to Amazon EMR on EKS

For any actions that you perform on Amazon EMR on EKS, you need a corresponding IAM permission for that action. You must create an IAM policy that allows you to perform the Amazon EMR on EKS actions and attach the policy to the IAM user or role that you use.

This topic provides steps for creating a new policy and attaching it to a user. It also covers the basic permissions that you need to set up your Amazon EMR on EKS environment. We recommend that you refine the permissions to specific resources whenever possible based on your business needs.

Creating a new IAM policy and attaching it to a user in the IAM console

Create a new IAM policy

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the left navigation pane of the IAM console, choose **Policies**.
3. On the **Policies** page, choose **Create Policy**.
4. In the **Create Policy** window, navigate to the **Edit JSON** tab. Create a policy document with one or more JSON statements as shown in the examples following this procedure. Next, choose **Review policy**.
5. On the **Review Policy** screen, enter your **Policy Name**, for example AmazonEMR0nEKSPolicy. Enter an optional description, and then choose **Create policy**.

Attach the policy to a user or role

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Policies**.
3. In the list of policies, select the check box next to the policy created in the previous section. You can use the **Filter** menu and the search box to filter the list of policies.
4. Choose **Policy actions**, and then choose **Attach**.
5. Choose the user or role to attach the policy to. You can use the **Filter** menu and the search box to filter the list of principal entities. After choosing the user or role to attach the policy to, choose **Attach policy**.

Permissions for managing virtual clusters

To manage virtual clusters in your AWS account, create an IAM policy with the following permissions. These permissions allow you to create, list, describe, and delete virtual clusters in your AWS account.

JSON

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

```

"Statement": [
  {
    "Effect": "Allow",
    "Action": [
      "iam:CreateServiceLinkedRole"
    ],
    "Resource": [
      "*"
    ],
    "Condition": {
      "StringLike": {
        "iam:AWSServiceName": "emr-containers.amazonaws.com"
      }
    },
    "Sid": "AllowIAMCreateservicelinkedrole"
  },
  {
    "Effect": "Allow",
    "Action": [
      "emr-containers:CreateVirtualCluster",
      "emr-containers:ListVirtualClusters",
      "emr-containers:DescribeVirtualCluster",
      "emr-containers>DeleteVirtualCluster"
    ],
    "Resource": [
      "*"
    ],
    "Sid": "AllowEMRCONTAINERSCreatevirtualcluster"
  }
]
}

```

Amazon EMR is integrated with Amazon EKS cluster access management (CAM), so you can automate configuration of the necessary AuthN and AuthZ policies to run Amazon EMR Spark jobs in namespaces of Amazon EKS clusters. To do so, you must have the following permissions:

```

{
  "Effect": "Allow",
  "Action": [
    "eks:CreateAccessEntry"
  ],

```

```

    "Resource":
      "arn:<AWS_PARTITION>:eks:<AWS_REGION>:<AWS_ACCOUNT_ID>:cluster/<EKS_CLUSTER_NAME>"
  },
  {
    "Effect": "Allow",
    "Action": [
      "eks:DescribeAccessEntry",
      "eks>DeleteAccessEntry",
      "eks>ListAssociatedAccessPolicies",
      "eks:AssociateAccessPolicy",
      "eks:DisassociateAccessPolicy"
    ],
    "Resource": "arn:<AWS_PARTITION>:eks:<AWS_REGION>:<AWS_ACCOUNT_ID>:access-
entry/<EKS_CLUSTER_NAME>/role/<AWS_ACCOUNT_ID>/AWSServiceRoleForAmazonEMRContainers/*"
  }
}

```

For more information, see [Automate enabling cluster access for Amazon EMR on EKS](#).

When the `CreateVirtualCluster` operation is invoked for the first time from an AWS account, you also need the `CreateServiceLinkedRole` permissions to create the service-linked role for Amazon EMR on EKS. For more information, see [Using service-linked roles for Amazon EMR on EKS](#).

Permissions for submitting jobs

To submit jobs on the virtual clusters in your AWS account, create an IAM policy with the following permissions. These permissions allow you to start, list, describe, and cancel job runs for the all virtual clusters in your account. You should consider adding permissions to list or describe virtual clusters, which allow you to check the state of the virtual cluster before submitting jobs.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "emr-containers:StartJobRun",
        "emr-containers:ListJobRuns",
        "emr-containers:DescribeJobRun",
        "emr-containers:CancelJobRun"
      ]
    }
  ]
}

```

```

    ],
    "Resource": [
        "*"
    ],
    "Sid": "AllowEMRCONTAINERSStartjobrun"
  }
]
}

```

Permissions for debugging and monitoring

To get access to logs pushed to Amazon S3 and CloudWatch, or to view application event logs in the Amazon EMR console, create an IAM policy with the following permissions. We recommend that you refine the permissions to specific resources whenever possible based on your business needs.

Important

If you haven't created an Amazon S3 bucket, you need to add `s3:CreateBucket` permission to the policy statement. If you haven't created a log group, you need to add `logs:CreateLogGroup` to the policy statement.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "emr-containers:DescribeJobRun",
        "elasticmapreduce:CreatePersistentAppUI",
        "elasticmapreduce:DescribePersistentAppUI",
        "elasticmapreduce:GetPersistentAppUIPresignedURL"
      ],
      "Resource": [
        "*"
      ],
      "Sid": "AllowEMRCONTAINERSDescribejobrun"
    }
  ]
}

```

```
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:ListBucket"
      ],
      "Resource": [
        "*"
      ],
      "Sid": "AllowS3GetObject"
    },
    {
      "Effect": "Allow",
      "Action": [
        "logs:Get*",
        "logs:DescribeLogGroups",
        "logs:DescribeLogStreams"
      ],
      "Resource": [
        "*"
      ],
      "Sid": "AllowLOGSGet"
    }
  ]
}
```

For more information about how to configure a job run to push logs to Amazon S3 and CloudWatch, see [Configure a job run to use S3 logs](#) and [Configure a job run to use CloudWatch Logs](#).

Register the Amazon EKS cluster with Amazon EMR

Registering your cluster is the final required step to set up Amazon EMR on EKS to run workloads.

Use the following command to create a virtual cluster with a name of your choice for the Amazon EKS cluster and namespace that you set up in previous steps.

Note

Each virtual cluster must have a unique name across all the EKS clusters. If two virtual clusters have the same name, the deployment process will fail even if the two virtual clusters belong to different EKS clusters.

```
aws emr-containers create-virtual-cluster \  
--name virtual_cluster_name \  
--container-provider '{  
  "id": "cluster_name",  
  "type": "EKS",  
  "info": {  
    "eksInfo": {  
      "namespace": "namespace_name"  
    }  
  }  
'
```

Alternatively, you can create a JSON file that includes the required parameters for the virtual cluster and then run the `create-virtual-cluster` command with the path to the JSON file. For more information, see [Managing virtual clusters](#).

Note

To validate the successful creation of a virtual cluster, view the status of virtual clusters using the `list-virtual-clusters` operation or by going to the **Virtual Clusters** page in the Amazon EMR console.

Submit a job run with StartJobRun

To submit a job run with a JSON file with specified parameters

1. Create a `start-job-run-request.json` file and specify the required parameters for your job run, as the following example JSON file demonstrates. For more information about the parameters, see [Options for configuring a job run](#).

```
{
```

```

"name": "myjob",
"virtualClusterId": "123456",
"executionRoleArn": "iam_role_name_for_job_execution",
"releaseLabel": "emr-6.2.0-latest",
"jobDriver": {
  "sparkSubmitJobDriver": {
    "entryPoint": "entryPoint_location",
    "entryPointArguments": ["argument1", "argument2", ...],
    "sparkSubmitParameters": "--class <main_class> --conf
spark.executor.instances=2 --conf spark.executor.memory=2G --conf
spark.executor.cores=2 --conf spark.driver.cores=1"
  }
},
"configurationOverrides": {
  "applicationConfiguration": [
    {
      "classification": "spark-defaults",
      "properties": {
        "spark.driver.memory": "2G"
      }
    }
  ],
  "monitoringConfiguration": {
    "persistentAppUI": "ENABLED",
    "cloudWatchMonitoringConfiguration": {
      "logGroupName": "my_log_group",
      "logStreamNamePrefix": "log_stream_prefix"
    },
    "s3MonitoringConfiguration": {
      "logUri": "s3://my_s3_log_location"
    }
  }
}
}

```

2. Use the `start-job-run` command with a path to the `start-job-run-request.json` file stored locally.

```

aws emr-containers start-job-run \
--cli-input-json file:///./start-job-run-request.json

```

To start a job run using the start-job-run command

1. Supply all the specified parameters in the StartJobRun command, as the following example demonstrates.

```
aws emr-containers start-job-run \
--virtual-cluster-id 123456 \
--name myjob \
--execution-role-arn execution-role-arn \
--release-label emr-6.2.0-latest \
--job-driver '{"sparkSubmitJobDriver": {"entryPoint": "entryPoint_location",
"entryPointArguments": ["argument1", "argument2", ...], "sparkSubmitParameters":
"--class <main_class> --conf spark.executor.instances=2 --conf
spark.executor.memory=2G --conf spark.executor.cores=2 --conf
spark.driver.cores=1"}}' \
--configuration-overrides '{"applicationConfiguration": [{"classification":
"spark-defaults", "properties": {"spark.driver.memory": "2G"}}],
"monitoringConfiguration": {"cloudWatchMonitoringConfiguration":
{"logGroupName": "log_group_name", "logStreamNamePrefix": "log_stream_prefix"},
"persistentAppUI": "ENABLED", "s3MonitoringConfiguration": {"logUri":
"s3://my_s3_log_location" }}}
```

2. For Spark SQL, supply all the specified parameters in the StartJobRun command, as the following example demonstrates.

```
aws emr-containers start-job-run \
--virtual-cluster-id 123456 \
--name myjob \
--execution-role-arn execution-role-arn \
--release-label emr-6.7.0-latest \
--job-driver '{"sparkSqlJobDriver": {"entryPoint": "entryPoint_location",
"sparkSqlParameters": "--conf spark.executor.instances=2 --conf
spark.executor.memory=2G --conf spark.executor.cores=2 --conf
spark.driver.cores=1"}}' \
--configuration-overrides '{"applicationConfiguration": [{"classification":
"spark-defaults", "properties": {"spark.driver.memory": "2G"}}],
"monitoringConfiguration": {"cloudWatchMonitoringConfiguration":
{"logGroupName": "log_group_name", "logStreamNamePrefix": "log_stream_prefix"},
"persistentAppUI": "ENABLED", "s3MonitoringConfiguration": {"logUri":
"s3://my_s3_log_location" }}}
```

Using job submitter classification

Overview

The Amazon EMR on EKS `StartJobRun` request creates a *job submitter* pod (also known as the *job-runner* pod) to spawn the Spark driver. You can use the `emr-job-submitter` classification to configure node selectors, add tolerations, customize logging, and make other modifications to the job submitter pod.

The following settings are available under the `emr-job-submitter` classification:

`jobsubmitter.node.selector.[selectorKey]`

Adds to the node selector of the job submitter pod, with key *selectorKey* and the value as the configuration value. For example, you can set `jobsubmitter.node.selector.identifier` to `myIdentifier` and the job submitter pod will have a node selector with a key `identifier` and a value `myIdentifier`. This can be used to specify which nodes the job submitter pod can be placed on. To add multiple node selector keys, set multiple configurations with this prefix.

`jobsubmitter.label.[labelKey]`

Adds to the labels of the job submitter pod, with key *labelKey* and the value as the configuration value. To add multiple labels, set multiple configurations with this prefix.

`jobsubmitter.annotation.[annotationKey]`

Adds to the annotations of the job submitter pod, with key *annotationKey* and the value as the configuration value. To add multiple annotations, set multiple configurations with this prefix.

`jobsubmitter.node.toleration.[tolerationKey]`

Adds [tolerations](#) to the job submitter pod. By default there are no tolerations added to the pod. The toleration's key will be *tolerationKey* and the toleration's value will be the configuration value. If the configuration value is set to a non-empty string, the operator will be `Equals`. If the configuration value is set to `""`, then the operator will be `Exists`.

`jobsubmitter.node.toleration.[tolerationKey].[effect]`

Adds a toleration effect to the prefixed *tolerationKey*. This field is required when adding tolerations. The allowed values for the effect field are `NoExecute`, `NoSchedule`, and `PreferNoSchedule`.

jobsubmitter.node.toleration.[*tolerationKey*].[*tolerationSeconds*]

Adds *tolerationSeconds* to the prefixed *tolerationKey*. Optional field. Only applicable when the effect is NoExecute.

jobsubmitter.scheduler.name

Sets a custom schedulerName for the job submitter pod.

jobsubmitter.logging

Enables or disables logging on the job submitter pod. When this is set to `DISABLED` the logging container is removed from the job submitter pod, which will disable any logging for this pod specified in the monitoringConfiguration, such as `s3MonitoringConfiguration` or `cloudWatchMonitoringConfiguration`. When this setting is not set or is set to any other value, logging on the job submitter pod is enabled.

jobsubmitter.logging.image

Sets a custom image to be used for the logging container on the job submitter pod.

jobsubmitter.logging.request.cores

Sets a custom value for the number of CPUs, in CPU units, for the logging container on the job submitter pod. By default, this is set to **100m**.

jobsubmitter.logging.request.memory

Sets a custom value for the amount of memory, in bytes, for the logging container on the job submitter pod. By default, this is set to **200Mi**. A mebibyte is a unit of measure that's similar to a megabyte.

jobsubmitter.container.image

Sets a custom image for the job submitter pod's job-runner container.

jobsubmitter.container.image.pullPolicy

Sets the [imagePullPolicy](#) for the job submitter pod's containers.

jobsubmitter.gracefulTermination

By default or when this config is set to `false`, after the job runner pod creates the driver pod it continues running, watches the driver pod, and periodically logs the driver status for the lifetime of the job. When this config is set to `true`, the job runner pod instead terminates immediately after starting the Spark driver pod. This means the job runner pod consumes CPU and memory for less time, but job runner logs are no longer available.

We recommend to place job submitter pods on On-Demand Instances. Placing job submitter pods on Spot instances might result in a job failure if the instance where the job submitter pod runs is subject to a Spot Instance interruption. You can also [place the job submitter pod in a single Availability Zone or use any Kubernetes labels that are applied to the nodes](#).

Job submitter classification examples

In this section

- [StartJobRun request with On-Demand node placement for the job submitter pod](#)
- [StartJobRun request with single-AZ node placement and Amazon EC2 instance type placement for the job submitter pod](#)
- [StartJobRun request with labels, annotations, and a custom scheduler for the job submitter pod](#)
- [StartJobRun request with a toleration applied to the job submitter pod with key dedicated, value graviton_machines, effect NoExecute, and a tolerationSeconds of 60 seconds](#)
- [StartJobRun request with logging disabled for the job submitter pod](#)
- [StartJobRun request with custom logging container image, CPU, and memory for the job submitter pod](#)
- [StartJobRun request with a custom job submitter container image and pull policy](#)
- [StartJobRun request with graceful termination enabled for the job submitter pod](#)

StartJobRun request with On-Demand node placement for the job submitter pod

```
cat >spark-python-in-s3-nodeselector-job-submitter.json << EOF
{
  "name": "spark-python-in-s3-nodeselector",
  "virtualClusterId": "virtual-cluster-id",
  "executionRoleArn": "execution-role-arn",
  "releaseLabel": "emr-6.11.0-latest",
  "jobDriver": {
    "sparkSubmitJobDriver": {
      "entryPoint": "s3://S3-prefix/trip-count.py",
      "sparkSubmitParameters": "--conf spark.driver.cores=5 --conf
spark.executor.memory=20G --conf spark.driver.memory=15G --conf
spark.executor.cores=6"
    }
  },
  "configurationOverrides": {
```

```

"applicationConfiguration": [
  {
    "classification": "spark-defaults",
    "properties": {
      "spark.dynamicAllocation.enabled": "false"
    }
  },
  {
    "classification": "emr-job-submitter",
    "properties": {
      "jobsubmitter.node.selector.eks.amazonaws.com/capacityType": "ON_DEMAND"
    }
  }
],
"monitoringConfiguration": {
  "cloudWatchMonitoringConfiguration": {
    "logGroupName": "/emr-containers/jobs",
    "logStreamNamePrefix": "demo"
  },
  "s3MonitoringConfiguration": {
    "logUri": "s3://joblogs"
  }
}
}
EOF
aws emr-containers start-job-run --cli-input-json file:///spark-python-in-s3-
nodeselector-job-submitter.json

```

StartJobRun request with single-AZ node placement and Amazon EC2 instance type placement for the job submitter pod

```

"configurationOverrides": {
  "applicationConfiguration": [
    {
      "classification": "emr-job-submitter",
      "properties": {
        "jobsubmitter.node.selector.topology.kubernetes.io/zone": "Availability Zone",
        "jobsubmitter.node.selector.node.kubernetes.io/instance-type": "m5.4xlarge"
      }
    }
  ]
}

```

StartJobRun request with labels, annotations, and a custom scheduler for the job submitter pod

```
"configurationOverrides": {
  "applicationConfiguration": [
    {
      "classification": "emr-job-submitter",
      "properties": {
        "jobsubmitter.label.label1": "value1",
        "jobsubmitter.label.label2": "value2",
        "jobsubmitter.annotation.ann1": "value1",
        "jobsubmitter.annotation.ann2": "value2",
        "jobsubmitter.scheduler.name": "custom-scheduler"
      }
    }
  ]
}
```

StartJobRun request with a toleration applied to the job submitter pod with key dedicated, value graviton_machines, effect NoExecute, and a tolerationSeconds of 60 seconds

```
"configurationOverrides": {
  "applicationConfiguration": [
    {
      "classification": "emr-job-submitter",
      "properties": {
        "jobsubmitter.node.toleration.dedicated": "graviton_machines",
        "jobsubmitter.node.toleration.dedicated.effect": "NoExecute",
        "jobsubmitter.node.toleration.dedicated.tolerationSeconds": "60"
      }
    }
  ]
}
```

StartJobRun request with logging disabled for the job submitter pod

```
"configurationOverrides": {
  "applicationConfiguration": [
    {
      "classification": "emr-job-submitter",
      "properties": {
        "jobsubmitter.logging": "DISABLED"
      }
    }
  ]
}
```



```

    }
  }
],
"monitoringConfiguration": {
  "cloudWatchMonitoringConfiguration": {
    "logGroupName": "/emr-containers/jobs",
    "logStreamNamePrefix": "demo"
  },
  "s3MonitoringConfiguration": {
    "logUri": "s3://joblogs"
  }
}
}
}

```

StartJobRun request with custom logging container image, CPU, and memory for the job submitter pod

```

"configurationOverrides": {
  "applicationConfiguration": [
    {
      "classification": "emr-job-submitter",
      "properties": {
        "jobsubmitter.logging.image": "YOUR_ECR_IMAGE_URL",
        "jobsubmitter.logging.request.memory": "200Mi",
        "jobsubmitter.logging.request.cores": "0.5"
      }
    }
  ],
  "monitoringConfiguration": {
    "cloudWatchMonitoringConfiguration": {
      "logGroupName": "/emr-containers/jobs",
      "logStreamNamePrefix": "demo"
    },
    "s3MonitoringConfiguration": {
      "logUri": "s3://joblogs"
    }
  }
}
}

```

StartJobRun request with a custom job submitter container image and pull policy

```

"configurationOverrides": {
  "applicationConfiguration": [

```

```
{
  "classification": "emr-job-submitter",
  "properties": {
    "jobsubmitter.container.image": "123456789012.dkr.ecr.us-west-2.amazonaws.com/emr6.11_custom_repo",
    "jobsubmitter.container.image.pullPolicy": "kubernetes pull policy"
  }
}
```

StartJobRun request with graceful termination enabled for the job submitter pod

```
"configurationOverrides": {
  "applicationConfiguration": [
    {
      "classification": "emr-job-submitter",
      "properties": {
        "jobsubmitter.gracefulTermination": "true"
      }
    }
  ]
}
```

Using Amazon EMR container defaults classification

Overview

The following settings are available under the `emr-containers-defaults` classification:

job-start-timeout

By default, a job will time out if it cannot start and it waits in the `SUBMITTED` state for 15 minutes. This configuration changes the number of seconds to wait before the job times out.

executor.logging

Enables or disables logging on the executor pods. When this is set to `DISABLED` the logging container is removed from the executor pods, which will disable any logging for these pods specified in the `monitoringConfiguration`, such as `s3MonitoringConfiguration` or `cloudWatchMonitoringConfiguration`. When this setting is not set or is set to any other value, logging on the executor pods is enabled.

logging.image

Sets a custom image to be used for the logging container on the driver and executor pods.

logging.request.cores

Sets a custom value for the number of CPUs, in CPU units, for the logging container on the driver and executor pods. By default, this is not set.

logging.request.memory

Sets a custom value for the amount of memory, in bytes, for the logging container on the driver and executor pods. By default, this is set to **512Mi**. A mebibyte is a unit of measure that's similar to a megabyte.

logging.eventLog.dir

Use this configuration to enable Persistent App UI and save Spark event logs to your own S3 location. Set `logging.eventLog.dir` to the S3 path for your event logs. In the `monitoringConfiguration`, set `persistentAppUI` to `ENABLED`. Don't set `spark.eventLog.dir` when you use `logging.eventLog.dir`; these two configurations are incompatible. If `spark.eventLog.dir` is set, that configuration takes priority and the logging container is unable to replicate Spark event logs. This means your S3 location specified by `logging.eventLog.dir` won't receive event logs and the Persistent App UI also won't work.

logging.nativeSidecar

When you set this property to `ENABLED` for Amazon EMR release 6.8.0 or higher, Amazon EMR configures the logging container on your Spark driver and executor pods as a Kubernetes native sidecar container (see details on the [Kubernetes website](#)) instead of a regular container. This means the logging container automatically restarts on failure and logging container failures won't cause the pod to fail.

Node version requirement

Your Amazon EKS nodes must be running Kubernetes version 1.29 or higher. Your EKS cluster version can differ from your node version. If your nodes are running a version lower than 1.29, Kubernetes does not enable the native sidecar feature and the logging container prevents the driver from starting, which leads to job timeouts.

Job submitter classification examples

In this section

- [StartJobRun request with custom job timeout](#)
- [StartJobRun request with logging disabled for executor pods](#)
- [StartJobRun request with custom logging container image, CPU, and memory for the driver and executor pods](#)
- [StartJobRun request with Spark event log Amazon S3 destination](#)
- [StartJobRun request with native sidecar logging](#)

StartJobRun request with custom job timeout

```
{
  "name": "spark-python",
  "virtualClusterId": "virtual-cluster-id",
  "executionRoleArn": "execution-role-arn",
  "releaseLabel": "emr-6.11.0-latest",
  "jobDriver": {
    "sparkSubmitJobDriver": {
      "entryPoint": "s3://S3-prefix/trip-count.py"
    }
  },
  "configurationOverrides": {
    "applicationConfiguration": [
      {
        "classification": "emr-containers-defaults",
        "properties": {
          "job-start-timeout": "1800"
        }
      }
    ],
    "monitoringConfiguration": {
      "cloudWatchMonitoringConfiguration": {
        "logGroupName": "/emr-containers/jobs",
        "logStreamNamePrefix": "demo"
      },
      "s3MonitoringConfiguration": {
        "logUri": "s3://joblogs"
      }
    }
  }
}
```

```
}
}
```

StartJobRun request with logging disabled for executor pods

```
"configurationOverrides": {
  "applicationConfiguration": [
    {
      "classification": "emr-containers-defaults",
      "properties": {
        "executor.logging": "DISABLED"
      }
    }
  ],
  "monitoringConfiguration": {
    "cloudWatchMonitoringConfiguration": {
      "logGroupName": "/emr-containers/jobs",
      "logStreamNamePrefix": "demo"
    },
    "s3MonitoringConfiguration": {
      "logUri": "s3://joblogs"
    }
  }
}
```

StartJobRun request with custom logging container image, CPU, and memory for the driver and executor pods

```
"configurationOverrides": {
  "applicationConfiguration": [
    {
      "classification": "emr-containers-defaults",
      "properties": {
        "logging.image": "YOUR_ECR_IMAGE_URL",
        "logging.request.memory": "200Mi",
        "logging.request.cores": "0.5"
      }
    }
  ],
  "monitoringConfiguration": {
    "cloudWatchMonitoringConfiguration": {
      "logGroupName": "/emr-containers/jobs",
      "logStreamNamePrefix": "demo"
    }
  }
}
```

```

    },
    "s3MonitoringConfiguration": {
      "logUri": "s3://joblogs"
    }
  }
}

```

Note

If the Fluentd logging container encounters an out-of-memory (OOM) error, increase the `logging.request.memory` value. For example, set it to `1Gi` to allocate more memory to the logging container and prevent OOM issues.

StartJobRun request with Spark event log Amazon S3 destination

The following example saves Spark event logs to your own Amazon S3 bucket while also enabling Persistent App UI. The `persistentAppUI` setting is `ENABLED` by default.

```

"configurationOverrides": {
  "applicationConfiguration": [
    {
      "classification": "emr-containers-defaults",
      "properties": {
        "logging.eventLog.dir": "s3://my-bucket/event-logs/"
      }
    }
  ],
  "monitoringConfiguration": {
    "persistentAppUI": "ENABLED"
  }
}

```

Note

Don't set `spark.eventLog.dir` in the `spark-defaults` classification when you use `logging.eventLog.dir`. These two configurations are incompatible. If `spark.eventLog.dir` is set, that configuration takes priority and the logging container is unable to replicate Spark event logs. This means your S3 location specified by

`logging.eventLog.dir` won't receive event logs and the Persistent App UI also won't work.

StartJobRun request with native sidecar logging

The following example enables native sidecar mode for the logging container on Spark driver and executor pods. When enabled, the logging container runs as a Kubernetes native sidecar that automatically restarts on failure and doesn't affect the state of your Spark pods.

```
"configurationOverrides": {
  "applicationConfiguration": [
    {
      "classification": "emr-containers-defaults",
      "properties": {
        "logging.nativeSidecar": "ENABLED"
      }
    }
  ],
  "monitoringConfiguration": {
    "s3MonitoringConfiguration": {
      "logUri": "s3://my-bucket/logs/"
    }
  }
}
```

Node version requirement

Your Amazon EKS nodes must be running Kubernetes version 1.29 or higher. Your EKS cluster version can differ from your node version. If your nodes are running a version lower than 1.29, Kubernetes does not enable the native sidecar feature and the logging container prevents the driver from starting, which leads to job timeouts.

Running Spark jobs with the Spark operator

Amazon EMR releases 6.10.0 and higher support the Kubernetes operator for Apache Spark, or *the Spark operator*, as a job submission model for Amazon EMR on EKS. With the Spark operator, you can deploy and manage Spark applications with the Amazon EMR release runtime on your own

Amazon EKS clusters. Once you deploy the Spark operator in your Amazon EKS cluster, you can directly submit Spark applications with the operator. The operator manages the lifecycle of Spark applications.

Note

Amazon EMR calculates pricing on Amazon EKS based on vCPU and memory consumption. This calculation applies to driver and executor pods. This calculation starts from when you download your Amazon EMR application image until the Amazon EKS pod terminates and is rounded to the nearest second. Pricing applies whenever pods run using an image derived from the Amazon EMR runtime, regardless of whether the image is stored in the public Amazon EMR Amazon ECR registry or your own private registry, and regardless of how the job is orchestrated (for example, the StartJobRun API, the Spark operator, or spark-submit).

Topics

- [Setting up the Spark operator for Amazon EMR on EKS](#)
- [Getting started with the Spark operator for Amazon EMR on EKS](#)
- [Use vertical autoscaling with the Spark operator for Amazon EMR on EKS](#)
- [Uninstalling the Spark operator for Amazon EMR on EKS](#)
- [Using monitoring configuration to monitor the Spark Kubernetes operator and Spark jobs](#)
- [Security and the Spark operator with Amazon EMR on EKS](#)

Setting up the Spark operator for Amazon EMR on EKS

Complete the following tasks to get set up before you install the Spark operator on Amazon EKS. If you've already signed up for Amazon Web Services (AWS) and have used Amazon EKS, you are almost ready to use Amazon EMR on EKS. Complete the following tasks to get set up for the Spark operator on Amazon EKS. If you've already completed any of the prerequisites, you can skip those and move on to the next one.

- [Install or update to the latest version of the AWS CLI](#) – If you've already installed the AWS CLI, confirm that you have the latest version.
- [Set up kubectl and eksctl](#) – eksctl is a command line tool that you use to communicate with Amazon EKS.

- [Install Helm](#) – The Helm package manager for Kubernetes helps you install and manage applications on your Kubernetes cluster.
- [Get started with Amazon EKS – eksctl](#) – Follow the steps to create a new Kubernetes cluster with nodes in Amazon EKS.
- [Select an Amazon EMR base image URI \(release 6.10.0 or higher\)](#) – the Spark operator is supported with Amazon EMR releases 6.10.0 and higher.

Getting started with the Spark operator for Amazon EMR on EKS

This topic helps you start to use the Spark operator on Amazon EKS by deploying a Spark application and a Schedule Spark application.

Install the Spark operator

Use the following steps to install the Kubernetes operator for Apache Spark.

1. If you haven't already, complete the steps in [Setting up the Spark operator for Amazon EMR on EKS](#).
2. Authenticate your Helm client to the Amazon ECR registry. In the following command, replace the *region-id* values with your preferred AWS Region, and the corresponding *ECR-registry-account* value for the Region from the [Amazon ECR registry accounts by Region](#) page.

```
aws ecr get-login-password \  
--region region-id | helm registry login \  
--username AWS \  
--password-stdin ECR-registry-account.dkr.ecr.region-id.amazonaws.com
```

3. Install the Spark operator with the following command.

For the Helm chart `--version` parameter, use your Amazon EMR release label with the `emr-` prefix and date suffix removed. For example, with the `emr-6.12.0-java17-latest` release, specify `6.12.0-java17`. The example in the following command uses the `emr-7.13.0-latest` release, so it specifies `7.13.0` for the Helm chart `--version`.

```
helm install spark-operator-demo \  
oci://895885662937.dkr.ecr.region-id.amazonaws.com/spark-operator \  
--set emrContainers.awsRegion=region-id \  

```

```
--version 7.13.0 \  
--namespace spark-operator \  
--create-namespace
```

By default, the command creates service account `emr-containers-sa-spark-operator` for the Spark operator. To use a different service account, provide the argument `serviceAccounts.sparkoperator.name`. For example:

```
--set serviceAccounts.sparkoperator.name my-service-account-for-spark-operator
```

If you want to [use vertical autoscaling with the Spark operator](#), add the following line to the installation command to allow webhooks for the operator:

```
--set webhook.enable=true
```

4. Verify that you installed the Helm chart with the `helm list` command:

```
helm list --namespace spark-operator -o yaml
```

The `helm list` command should return your newly-deployed Helm chart release information:

```
app_version: v1beta2-1.3.8-3.1.1  
chart: spark-operator-7.13.0  
name: spark-operator-demo  
namespace: spark-operator  
revision: "1"  
status: deployed  
updated: 2023-03-14 18:20:02.721638196 +0000 UTC
```

5. Complete installation with any additional options that you require. For more information, see the [spark-on-k8s-operator](#) documentation on GitHub.

Run a Spark application

The Spark operator is supported with Amazon EMR 6.10.0 or higher. When you install the Spark operator, it creates the service account `emr-containers-sa-spark` to run Spark applications by default. Use the following steps to run a Spark application with the Spark operator on Amazon EMR on EKS 6.10.0 or higher.

1. Before you can run a Spark application with the Spark operator, complete the steps in [Setting up the Spark operator for Amazon EMR on EKS](#) and [Install the Spark operator](#).
2. Create a SparkApplication definition file `spark-pi.yaml` with the following example contents:

```
apiVersion: "sparkoperator.k8s.io/v1beta2"
kind: SparkApplication
metadata:
  name: spark-pi
  namespace: spark-operator
spec:
  type: Scala
  mode: cluster
  image: "895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/emr-6.10.0:latest"
  imagePullPolicy: Always
  mainClass: org.apache.spark.examples.SparkPi
  mainApplicationFile: "local:///usr/lib/spark/examples/jars/spark-examples.jar"
  sparkVersion: "3.3.1"
  restartPolicy:
    type: Never
  volumes:
    - name: "test-volume"
      hostPath:
        path: "/tmp"
        type: Directory
  driver:
    cores: 1
    coreLimit: "1200m"
    memory: "512m"
    labels:
      version: 3.3.1
    serviceAccount: emr-containers-sa-spark
    volumeMounts:
      - name: "test-volume"
        mountPath: "/tmp"
  executor:
    cores: 1
    instances: 1
    memory: "512m"
    labels:
      version: 3.3.1
    volumeMounts:
      - name: "test-volume"
```

```
mountPath: "/tmp"
```

3. Now, submit the Spark application with the following command. This will also create a SparkApplication object named spark-pi:

```
kubectl apply -f spark-pi.yaml
```

4. Check events for the SparkApplication object with the following command:

```
kubectl describe sparkapplication spark-pi --namespace spark-operator
```

For more information on submitting applications to Spark through the Spark operator, see [Using a SparkApplication](#) in the spark-on-k8s-operator documentation on GitHub.

Use Amazon S3 for storage

To use Amazon S3 as your file storage option, add the following configurations to your YAML file.

```
hadoopConf:
# EMRFS filesystem
  fs.s3.customAWSCredentialsProvider:
com.amazonaws.auth.WebIdentityTokenCredentialsProvider
  fs.s3.impl: com.amazon.ws.emr.hadoop.fs.EmrFileSystem
  fs.AbstractFileSystem.s3.impl: org.apache.hadoop.fs.s3.EMRFSDelegate
  fs.s3.buffer.dir: /mnt/s3
  fs.s3.getObject.initialSocketTimeoutMilliseconds: "2000"
  mapreduce.fileoutputcommitter.algorithm.version.emr_internal_use_only.EmrFileSystem:
"2"
  mapreduce.fileoutputcommitter.cleanup-
failures.ignored.emr_internal_use_only.EmrFileSystem: "true"
sparkConf:
# Required for EMR Runtime
  spark.driver.extraClassPath: /usr/lib/hadoop-lzo/lib/*:/usr/lib/hadoop/hadoop-
aws.jar:/usr/share/aws/aws-java-sdk/*:/usr/share/aws/emr/emrfs/conf:/usr/share/aws/
emr/emrfs/lib/*:/usr/share/aws/emr/emrfs/auxlib/*:/usr/share/aws/emr/security/conf:/
usr/share/aws/emr/security/lib/*:/usr/share/aws/hmclient/lib/aws-glue-datacatalog-
spark-client.jar:/usr/share/java/Hive-JSON-Serde/hive-openx-serde.jar:/usr/share/aws/
sagemaker-spark-sdk/lib/sagemaker-spark-sdk.jar:/home/hadoop/extrajars/*
  spark.driver.extraLibraryPath: /usr/lib/hadoop/lib/native:/usr/lib/hadoop-lzo/lib/
native:/docker/usr/lib/hadoop/lib/native:/docker/usr/lib/hadoop-lzo/lib/native
  spark.executor.extraClassPath: /usr/lib/hadoop-lzo/lib/*:/usr/lib/hadoop/hadoop-
aws.jar:/usr/share/aws/aws-java-sdk/*:/usr/share/aws/emr/emrfs/conf:/usr/share/aws/
```

```
emr/emrfs/lib/*:/usr/share/aws/emr/emrfs/auxlib/*:/usr/share/aws/emr/security/conf:/usr/share/aws/emr/security/lib/*:/usr/share/aws/hmclient/lib/aws-glue-datacatalog-spark-client.jar:/usr/share/java/Hive-JSON-Serde/hive-openx-serde.jar:/usr/share/aws/sagemaker-spark-sdk/lib/sagemaker-spark-sdk.jar:/home/hadoop/extrajars/*  
spark.executor.extraLibraryPath: /usr/lib/hadoop/lib/native:/usr/lib/hadoop-lzo/lib/native:/docker/usr/lib/hadoop/lib/native:/docker/usr/lib/hadoop-lzo/lib/native
```

If you use Amazon EMR releases 7.2.0 and higher, the configurations are included by default. In that case, you can set the file path to `s3://<bucket_name>/<file_path>` instead of `local://<file_path>` in the Spark application YAML file.

Then submit the Spark application as normal.

Use vertical autoscaling with the Spark operator for Amazon EMR on EKS

Starting with Amazon EMR 7.0, you can use Amazon EMR on EKS vertical autoscaling to simplify resource management. It automatically tunes memory and CPU resources to adapt to the needs of the workload that you provide for Amazon EMR Spark applications. For more information, see [Using vertical autoscaling with Amazon EMR Spark jobs](#).

This section describes how to configure the Spark operator to use vertical autoscaling.

Prerequisites

Before you configure monitoring, be sure to complete the following setup tasks:

- Complete the steps in [Setting up the Spark operator for Amazon EMR on EKS](#).
- (optional) If you previously installed an older version of the Spark operator, delete the SparkApplication/ScheduledSparkApplication CRD.

```
kubectl delete crd sparkApplication  
kubectl delete crd scheduledSparkApplication
```

- Complete the steps in [Install the Spark operator](#). In step 3, add the following line to the installation command to allow webhooks for the operator:

```
--set webhook.enable=true
```

- Complete the steps in [Setting up vertical autoscaling for Amazon EMR on EKS](#).
- Give access to the files in your Amazon S3 location:

1. Annotate your driver and operator service account with the `JobExecutionRole` that has S3 permissions.

```
kubectl annotate serviceaccount -n spark-operator emr-containers-sa-spark
eks.amazonaws.com/role-arn=JobExecutionRole
kubectl annotate serviceaccount -n spark-operator emr-containers-sa-spark-
operator eks.amazonaws.com/role-arn=JobExecutionRole
```

2. Update the trust policy of your job execution role in that namespace.

```
aws emr-containers update-role-trust-policy \
--cluster-name cluster \
--namespace ${Namespace} \
--role-name iam_role_name_for_job_execution
```

3. Edit the IAM role trust policy of your job execution role and update the serviceaccount from `emr-containers-sa-spark-*-*xxxx` to `emr-containers-sa-*`.

```
{
  "Effect": "Allow",
  "Principal": {
    "Federated": "OIDC-provider"
  },
  "Action": "sts:AssumeRoleWithWebIdentity",
  "Condition": {
    "StringLike": {
      "OIDC": "system:serviceaccount:${Namespace}:emr-containers-sa-*"
    }
  }
}
```

4. If you're using Amazon S3 as your file storage, add the following defaults to your yaml file.

```
hadoopConf:
# EMRFS filesystem
  fs.s3.customAWSCredentialsProvider:
    com.amazonaws.auth.WebIdentityTokenCredentialsProvider
  fs.s3.impl: com.amazon.ws.emr.hadoop.fs.EmrFileSystem
  fs.AbstractFileSystem.s3.impl: org.apache.hadoop.fs.s3.EMRFSDelegate
  fs.s3.buffer.dir: /mnt/s3
  fs.s3.getObject.initialSocketTimeoutMilliseconds: "2000"
```

```
mapreduce.fileoutputcommitter.algorithm.version.emr_internal_use_only.EmrFileSystem:
"2"
mapreduce.fileoutputcommitter.cleanup-
failures.ignored.emr_internal_use_only.EmrFileSystem: "true"
sparkConf:
# Required for EMR Runtime
spark.driver.extraClassPath: /usr/lib/hadoop-lzo/lib/*:/usr/lib/hadoop/hadoop-
aws.jar:/usr/share/aws/aws-java-sdk/*:/usr/share/aws/emr/emrfs/conf:/usr/share/
aws/emr/emrfs/lib/*:/usr/share/aws/emr/emrfs/auxlib/*:/usr/share/aws/emr/
security/conf:/usr/share/aws/emr/security/lib/*:/usr/share/aws/hmclient/lib/aws-
glue-datacatalog-spark-client.jar:/usr/share/java/Hive-JSON-Serde/hive-openx-
serde.jar:/usr/share/aws/sagemaker-spark-sdk/lib/sagemaker-spark-sdk.jar:/home/
hadoop/extrajars/*
spark.driver.extraLibraryPath: /usr/lib/hadoop/lib/native:/usr/lib/hadoop-lzo/
lib/native:/docker/usr/lib/hadoop/lib/native:/docker/usr/lib/hadoop-lzo/lib/
native
spark.executor.extraClassPath: /usr/lib/hadoop-lzo/lib/*:/usr/lib/hadoop/
hadoop-aws.jar:/usr/share/aws/aws-java-sdk/*:/usr/share/aws/emr/emrfs/conf:/usr/
share/aws/emr/emrfs/lib/*:/usr/share/aws/emr/emrfs/auxlib/*:/usr/share/aws/emr/
security/conf:/usr/share/aws/emr/security/lib/*:/usr/share/aws/hmclient/lib/aws-
glue-datacatalog-spark-client.jar:/usr/share/java/Hive-JSON-Serde/hive-openx-
serde.jar:/usr/share/aws/sagemaker-spark-sdk/lib/sagemaker-spark-sdk.jar:/home/
hadoop/extrajars/*
spark.executor.extraLibraryPath: /usr/lib/hadoop/lib/native:/usr/lib/hadoop-
lzo/lib/native:/docker/usr/lib/hadoop/lib/native:/docker/usr/lib/hadoop-lzo/lib/
native
```

Run a job with vertical autoscaling on the Spark operator

Before you can run a Spark application with the Spark operator, you must complete the steps in [Prerequisites](#).

To use vertical autoscaling with the Spark operator, add the following configuration to the driver for your Spark Application spec to turn on vertical autoscaling:

```
dynamicSizing:
  mode: Off
  signature: "my-signature"
```

This configuration enables vertical autoscaling and is a required signature configuration that lets you choose a signature for your job.

For more information on the configurations and parameter values, see [Configuring vertical autoscaling for Amazon EMR on EKS](#). By default, your job submits in the monitoring-only **Off** mode of vertical autoscaling. This monitoring state lets you compute and view resource recommendations without performing autoscaling. For more information, see [Vertical autoscaling modes](#).

The following is a sample SparkApplication definition file named `spark-pi.yaml` with the required configurations to use vertical autoscaling.

```
apiVersion: "sparkoperator.k8s.io/v1beta2"
kind: SparkApplication
metadata:
  name: spark-pi
  namespace: spark-operator
spec:
  type: Scala
  mode: cluster
  image: "895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/emr-7.13.0:latest"
  imagePullPolicy: Always
  mainClass: org.apache.spark.examples.SparkPi
  mainApplicationFile: "local:///usr/lib/spark/examples/jars/spark-examples.jar"
  sparkVersion: "3.4.1"
  dynamicSizing:
    mode: Off
    signature: "my-signature"
  restartPolicy:
    type: Never
  volumes:
    - name: "test-volume"
      hostPath:
        path: "/tmp"
        type: Directory
  driver:
    cores: 1
    coreLimit: "1200m"
    memory: "512m"
    labels:
      version: 3.4.1
    serviceAccount: emr-containers-sa-spark
    volumeMounts:
      - name: "test-volume"
        mountPath: "/tmp"
  executor:
```



```
cores: 1
instances: 1
memory: "512m"
labels:
  version: 3.4.1
volumeMounts:
  - name: "test-volume"
    mountPath: "/tmp"
```

Now, submit the Spark application with the following command. This will also create a SparkApplication object named spark-pi:

```
kubectl apply -f spark-pi.yaml
```

For more information on submitting applications to Spark through the Spark operator, see [Using a SparkApplication](#) in the spark-on-k8s-operator documentation on GitHub.

Verifying the vertical autoscaling functionality

To verify that vertical autoscaling works correctly for the submitted job, use kubectl to get the verticalpodautoscaler custom resource and view your scaling recommendations.

```
kubectl get verticalpodautoscalers --all-namespaces \
-l=emr-containers.amazonaws.com/dynamic.sizing.signature=my-signature
```

The output from this query should resemble the following:

NAMESPACE		NAME			MODE
CPU	MEM	PROVIDED	AGE		
spark-operator	580026651	ds-p73j6mkosvc4xeb3gr7x4xol2bfcw5evqimzqojrlsvj3giozuq-vpa	True	15m	Off

If your output doesn't look similar or contains an error code, see [Troubleshooting Amazon EMR on EKS vertical autoscaling](#) for steps to help resolve the issue.

To remove the pods and applications, run the following command:

```
kubectl delete sparkapplication spark-pi
```

Uninstalling the Spark operator for Amazon EMR on EKS

Use the following steps to uninstall the Spark operator.

1. Delete the Spark operator using the correct namespace. For this example, the namespace is `spark-operator-demo`.

```
helm uninstall spark-operator-demo -n spark-operator
```

2. Delete the Spark operator service account:

```
kubectl delete sa emr-containers-sa-spark-operator -n spark-operator
```

3. Delete the Spark operator CustomResourceDefinitions (CRDs):

```
kubectl delete crd sparkapplications.sparkoperator.k8s.io  
kubectl delete crd scheduledsparkapplications.sparkoperator.k8s.io
```

Using monitoring configuration to monitor the Spark Kubernetes operator and Spark jobs

Monitoring configuration lets you easily set up log archiving of your Spark application and operator logs to Amazon S3 or to Amazon CloudWatch. You can choose one or both. Doing so adds a log agent sidecar to your spark operator pod, driver, and executor pods, and subsequently forwards these components' logs to your configured sinks.

Prerequisites

Before you configure monitoring, be sure to complete the following setup tasks:

1. (Optional) If you previously installed an older version of the Spark operator, delete the *SparkApplication/ScheduledSparkApplication* CRD.

```
kubectl delete crd scheduledsparkapplications.sparkoperator.k8s.io  
kubectl delete crd sparkapplications.sparkoperator.k8s.io
```

2. Create an operator/job execution role in IAM if you don't have one already.
3. Run the following command to update the trust policy of the operator/job execution role you just created:

```
aws emr-containers update-role-trust-policy \
--cluster-name cluster \
--namespace namespace \
--role-name iam_role_name_for_operator/job_execution_role
```

4. Edit the IAM role trust policy of your operator/job execution role to the following:

```
{
  "Effect": "Allow",
  "Principal": {
    "Federated": "${OIDC-provider}"
  },
  "Action": "sts:AssumeRoleWithWebIdentity",
  "Condition": {
    "StringLike": {
      "OIDC_PROVIDER:sub": "system:serviceaccount:${Namespace}:emr-
containers-sa-*"
    }
  }
}
```

5. Create a *monitoringConfiguration* policy in IAM with following permissions:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "logs:DescribeLogStreams",
        "logs:CreateLogStream",
        "logs:CreateLogGroup",
        "logs:PutLogEvents"
      ],
      "Resource": [
        "arn:aws:logs::*:log-group:log_group_name",
        "arn:aws:logs::*:log-group:log_group_name:"
      ],
      "Sid": "AllowLOGSDescribelogstreams"
    }
  ]
}
```

```

    },
    {
      "Effect": "Allow",
      "Action": [
        "logs:DescribeLogGroups"
      ],
      "Resource": [
        "*"
      ],
      "Sid": "AllowLOGSDescribeloggroups"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::bucket_name",
        "arn:aws:s3:::bucket_name/*"
      ],
      "Sid": "AllowS3Putobject"
    }
  ]
}

```

6. Attach the above policy to your operator/job execution role.

Spark Operator Logs

You can define monitoring configuration in the following way when doing `helm install`:

```

helm install spark-operator spark-operator \
--namespace namespace \
--set emrContainers.awsRegion=aws_region \
--set emrContainers.monitoringConfiguration.image=log_agent_image_url \
--set
  emrContainers.monitoringConfiguration.s3MonitoringConfiguration.logUri=S3_bucket_uri \
--set
  emrContainers.monitoringConfiguration.cloudWatchMonitoringConfiguration.logGroupName=log_group
\

```

```
--set
  emrContainers.monitoringConfiguration.cloudWatchMonitoringConfiguration.logStreamNamePrefix=Lo
\
--set emrContainers.monitoringConfiguration.sideCarResources.limits.cpuLimit=500m \
--set emrContainers.monitoringConfiguration.sideCarResources.limits.memoryLimit=512Mi \
--set
  emrContainers.monitoringConfiguration.containerLogRotationConfiguration.rotationSize=2GB
\
--set
  emrContainers.monitoringConfiguration.containerLogRotationConfiguration.maxFilesToKeep=10
\
--set webhook.enable=true \
--set emrContainers.operatorExecutionRoleArn=operator_execution_role_arn
```

Monitoring configuration

The following are the available configuration options under **monitoringConfiguration**.

- **Image** (optional) – Log agent image url. Will fetch by `emrReleaseLabel` if not provided.
- **s3MonitoringConfiguration** – Set this option to archive to Amazon S3.
 - **logUri** – (required) – The Amazon S3 bucket path where you want to store your logs.
 - The following are sample formats for the Amazon S3 bucket paths, after the logs are uploaded. The first example shows no log rotation enabled.

```
s3://${logUri}/${POD_NAME}/operator/stdout.gz
s3://${logUri}/${POD_NAME}/operator/stderr.gz
```

Log rotation enabled by default. You can see both a rotated file, with an incrementing index, and a current file, which is the same as the previous sample.

```
s3://${logUri}/${POD_NAME}/operator/stdout_YYYYMMDD_index.gz
s3://${logUri}/${POD_NAME}/operator/stderr_YYYYMMDD_index.gz
```

- **cloudWatchMonitoringConfiguration** – The configuration key to set up forwarding to Amazon CloudWatch.
 - **logGroupName** (required) – Name of the Amazon CloudWatch log group that you want to send logs to. The group automatically gets created if it doesn't exist.
 - **logStreamNamePrefix** (optional) – Name of the log stream that you want to send logs into. The default value is an empty string. The format in Amazon CloudWatch is as follows:

```
`${logStreamNamePrefix}/${POD_NAME}/STDOUT or STDERR
```

- **sideCarResources** (optional) – The configuration key to set resource limits on the launched Fluentd sidecar container.
 - **memoryLimit** (optional) – The memory limit. Adjust according to your needs. The default is 512Mi.
 - **cpuLimit** (optional) – The CPU limit. Adjust according to your needs. The default is 500m.
- **containerLogRotationConfiguration** (optional) – Controls the container log rotation behavior. It is enabled by default.
 - **rotationSize** (required) – Specifies file size for the log rotation. The range of possible values is from 2KB to 2GB. The numeric unit portion of the rotationSize parameter is passed as an integer. Since decimal values aren't supported, you can specify a rotation size of 1.5GB, for example, with the value 1500MB. The default is 2GB.
 - **maxFilesToKeep** (required) – Specifies the maximum number of files to retain in the container after rotation has taken place. The minimum value is 1, and the maximum value is 50. The default is 10.

After configured *monitoringConfiguration*, you should be able to check spark operator pod logs on an Amazon S3 bucket or Amazon CloudWatch or both. For an Amazon S3 bucket, you need to wait 2 minutes for the first log file to get flushed.

To find the logs in Amazon CloudWatch, you can navigate to the following: **CloudWatch > Log groups > *Log group name* > *Pod name*/operator/stderr**

Or you can navigate to: **CloudWatch > Log groups > *Log group name* > *Pod name*/operator/stdout**

Spark Application Logs

You can define this configuration in the following way.

```
apiVersion: "sparkoperator.k8s.io/v1beta2"
kind: SparkApplication
metadata:
  name: spark-pi
  namespace: namespace
spec:
  type: Scala
```

```
mode: cluster
imagePullPolicy: Always
mainClass: org.apache.spark.examples.SparkPi
mainApplicationFile: "local:///usr/lib/spark/examples/jars/spark-examples.jar"
sparkVersion: "3.3.1"
emrReleaseLabel: emr_release_label
executionRoleArn: job_execution_role_arn
restartPolicy:
  type: Never
volumes:
  - name: "test-volume"
    hostPath:
      path: "/tmp"
      type: Directory
driver:
  cores: 1
  coreLimit: "1200m"
  memory: "512m"
  labels:
    version: 3.3.1
  volumeMounts:
    - name: "test-volume"
      mountPath: "/tmp"
executor:
  cores: 1
  instances: 1
  memory: "512m"
  labels:
    version: 3.3.1
  volumeMounts:
    - name: "test-volume"
      mountPath: "/tmp"
monitoringConfiguration:
  image: "log_agent_image"
  s3MonitoringConfiguration:
    logUri: "S3_bucket_uri"
  cloudWatchMonitoringConfiguration:
    logGroupName: "log_group_name"
    logStreamNamePrefix: "log_stream_prefix"
sidecarResources:
  limits:
    cpuLimit: "500m"
    memoryLimit: "250Mi"
containerLogRotationConfiguration:
```

```
rotationSize: "2GB"
maxFilesToKeep: "10"
```

The following are the available configuration options under **monitoringConfiguration**.

- **Image** (optional) – Log agent image url. Will fetch by `emrReleaseLabel` if not provided.
- **s3MonitoringConfiguration** – Set this option to archive to Amazon S3.
 - **logUri** (required) – The Amazon S3 bucket path where you want to store your logs. The first example shows no log rotation enabled:

```
s3://${logUri}/${APPLICATION_NAME}-${APPLICATION_UID}/${POD_NAME}/stdout.gz
s3://${logUri}/${APPLICATION_NAME}-${APPLICATION_UID}/${POD_NAME}/stderr.gz
```

Log rotation is enabled by default. You can use both a rotated file (with incrementing index) and a current file (one without the date stamp).

```
s3://${logUri}/${APPLICATION_NAME}-${APPLICATION_UID}/${POD_NAME}/
stdout_YYYYMMDD_index.gz
s3://${logUri}/${APPLICATION_NAME}-${APPLICATION_UID}/${POD_NAME}/
stderr_YYYYMMDD_index.gz
```

- **cloudWatchMonitoringConfiguration** – The configuration key to set up forwarding to Amazon CloudWatch.
 - **logGroupName** (required) – The name of the Cloudwatch log group that you want to send logs to. The group automatically is created if it doesn't exist.
 - **logStreamNamePrefix** (optional) – The Name of the log stream that you want to send logs into. The default value is an empty string. The format in CloudWatch is as follows:

```
${logStreamNamePrefix}/${APPLICATION_NAME}-${APPLICATION_UID}/${POD_NAME}/stdout
${logStreamNamePrefix}/${APPLICATION_NAME}-${APPLICATION_UID}/${POD_NAME}/stderr
```

- **sideCarResources** (optional) – The configuration key to set resource limits on the launched Fluentd sidecar container.
 - **memoryLimit** (optional) – The memory limit. Adjust according to your needs. The default is 250Mi.
 - **cpuLimit** – The CPU limit. Adjust according to your needs. The default is 500m.
- **containerLogRotationConfiguration** (optional) – Controls the container log rotation behavior. It is enabled by default.

- **rotationSize** (required) – Specifies file size for the log rotation. The range of possible values is from 2KB to 2GB. The numeric unit portion of the rotationSize parameter is passed as an integer. Since decimal values aren't supported, you can specify a rotation size of 1.5GB, for example, with the value 1500MB. The default is 2GB.
- **maxFilesToKeep** (required) – Specifies the maximum number of files to retain in the container after rotation has taken place. The minimum value is 1. The maximum value is 50. The default is 10.

After configuring monitoringConfiguration, you should be able to check your spark application driver and executor logs on an Amazon S3 bucket or CloudWatch or both. For an Amazon S3 bucket, you need to wait 2 minutes for the first log file to be flushed. For example, in Amazon S3, the bucket path appears like the following:

Amazon S3 > **Buckets** > *Bucket name* > *Spark application name - UUID* > *Pod Name* > **stderr.gz**

Or:

Amazon S3 > **Buckets** > *Bucket name* > *Spark application name - UUID* > *Pod Name* > **stdout.gz**

In CloudWatch, the path appears like the following:

CloudWatch > **Log groups** > *Log group name* > *Spark application name - UUID/ Pod name* / **stderr**

Or:

CloudWatch > **Log groups** > *Log group name* > *Spark application name - UUID/ Pod name* / **stdout**

Security and the Spark operator with Amazon EMR on EKS

There are a couple ways to set up cluster-access permissions when you use the Spark operator. The first is to use role-based access control, Role-based access control (RBAC) restricts access based on a person's role within an organization. It has become a primary way to handle access. The second access method is to assume an AWS Identity and Access Management role, which provides resource access by means of specific assigned permissions.

Topics

- [Setting up cluster access permissions with role-based access control \(RBAC\)](#)
- [Setting up cluster access permissions with IAM roles for service accounts \(IRSA\)](#)

Setting up cluster access permissions with role-based access control (RBAC)

To deploy the Spark operator, Amazon EMR on EKS creates two roles and service accounts for the Spark operator and the Spark apps.

Topics

- [Operator service account and role](#)
- [Spark service account and role](#)

Operator service account and role

Amazon EMR on EKS creates the **operator service account and role** to manage SparkApplications for Spark jobs and for other resources such as services.

The default name for this service account is `emr-containers-sa-spark-operator`.

The following rules apply to this service role:

```
rules:
- apiGroups:
  - ""
  resources:
  - pods
  verbs:
  - "*"
- apiGroups:
  - ""
  resources:
  - services
  - configmaps
  - secrets
  verbs:
  - create
  - get
  - delete
  - update
- apiGroups:
  - extensions
```

```
- networking.k8s.io
resources:
- ingresses
verbs:
- create
- get
- delete
- apiGroups:
  - ""
resources:
- nodes
verbs:
- get
- apiGroups:
  - ""
resources:
- events
verbs:
- create
- update
- patch
- apiGroups:
  - ""
resources:
- resourcequotas
verbs:
- get
- list
- watch
- apiGroups:
  - apiextensions.k8s.io
resources:
- customresourcedefinitions
verbs:
- create
- get
- update
- delete
- apiGroups:
  - admissionregistration.k8s.io
resources:
- mutatingwebhookconfigurations
- validatingwebhookconfigurations
verbs:
```

```

- create
- get
- update
- delete
- apiGroups:
  - sparkoperator.k8s.io
  resources:
  - sparkapplications
  - sparkapplications/status
  - scheduledsparkapplications
  - scheduledsparkapplications/status
  verbs:
  - "*"
  {{- if .Values.batchScheduler.enable }}
  # required for the `volcano` batch scheduler
- apiGroups:
  - scheduling.incubator.k8s.io
  - scheduling.sigs.dev
  - scheduling.volcano.sh
  resources:
  - podgroups
  verbs:
  - "*"
  {{- end }}
  {{ if .Values.webhook.enable }}
- apiGroups:
  - batch
  resources:
  - jobs
  verbs:
  - delete
  {{- end }}

```

Spark service account and role

A Spark driver pod needs a Kubernetes service account in the same namespace as the pod. This service account needs permissions to create, get, list, patch and delete executor pods, and to create a Kubernetes headless service for the driver. The driver fails and exits without the service account unless the default service account in the pod's namespace has the required permissions.

The default name for this service account is `emr-containers-sa-spark`.

The following rules apply to this service role:

```
rules:
- apiGroups:
  - ""
  resources:
  - pods
  verbs:
  - "*"
- apiGroups:
  - ""
  resources:
  - services
  verbs:
  - "*"
- apiGroups:
  - ""
  resources:
  - configmaps
  verbs:
  - "*"
- apiGroups:
  - ""
  resources:
  - persistentvolumeclaims
  verbs:
  - "*"
```

Setting up cluster access permissions with IAM roles for service accounts (IRSA)

This section uses an example to demonstrate how to configure a Kubernetes service account to assume an AWS Identity and Access Management role. Pods that use the service account can then access any AWS service that the role has permissions to access.

The following example runs a Spark application to count the words from a file in Amazon S3. To do this, you can set up IAM roles for service accounts (IRSA) to authenticate and authorize Kubernetes service accounts.

Note

This example uses the "spark-operator" namespace for the Spark operator and for the namespace where you submit the Spark application.

Prerequisites

Before you try the example on this page, complete the following prerequisites:

- [Get set up for the Spark operator.](#)
- [Install the Spark operator.](#)
- [Create an Amazon S3 bucket.](#)
- Save your favorite poem in a text file named `poem.txt`, and upload the file to your S3 bucket. The Spark application that you create on this page will read the contents of the text file. For more information on uploading files to S3, see [Upload an object to your bucket](#) in the *Amazon Simple Storage Service User Guide*.

Configure a Kubernetes service account to assume an IAM role

Use the following steps to configure a Kubernetes service account to assume an IAM role that pods can use to access AWS services that the role has permissions to access.

1. After completing the [Prerequisites](#), use the AWS Command Line Interface to create an `example-policy.json` file that allows read-only access to the file that you uploaded to Amazon S3:

```
cat >example-policy.json <<EOF
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::my-pod-bucket",
        "arn:aws:s3:::my-pod-bucket/*"
      ]
    }
  ]
}
EOF
```

2. Then, create an IAM policy example-policy:

```
aws iam create-policy --policy-name example-policy --policy-document file:///
example-policy.json
```

3. Next, create an IAM role example-role and associate it with a Kubernetes service account for the Spark driver:

```
eksctl create iamserviceaccount --name driver-account-sa --namespace spark-operator \
--cluster my-cluster --role-name "example-role" \
--attach-policy-arn arn:aws:iam::111122223333:policy/example-policy --approve
```

4. Create a yaml file with the cluster role bindings that are required for the Spark driver service account:

```
cat >spark-rbac.yaml <<EOF
apiVersion: v1
kind: ServiceAccount
metadata:
  name: driver-account-sa
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: spark-role
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: edit
subjects:
- kind: ServiceAccount
  name: driver-account-sa
  namespace: spark-operator
EOF
```

5. Apply the cluster role binding configurations:

```
kubectl apply -f spark-rbac.yaml
```

The kubectl command should confirm successful creation of the account:

```
serviceaccount/driver-account-sa created
clusterrolebinding.rbac.authorization.k8s.io/spark-role configured
```

Running an application from the Spark operator

After you [configure the Kubernetes service account](#), you can run a Spark application that counts the number of words in the text file that you uploaded as part of the [Prerequisites](#).

1. Create a new file `word-count.yaml`, with a `SparkApplication` definition for your word-count application, based on Amazon EMR version 6.

```
cat >word-count.yaml <<EOF
apiVersion: "sparkoperator.k8s.io/v1beta2"
kind: SparkApplication
metadata:
  name: word-count
  namespace: spark-operator
spec:
  type: Java
  mode: cluster
  image: "895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/emr-6.10.0:latest"
  imagePullPolicy: Always
  mainClass: org.apache.spark.examples.JavaWordCount
  mainApplicationFile: local:///usr/lib/spark/examples/jars/spark-examples.jar
  arguments:
    - s3://my-pod-bucket/poem.txt
  hadoopConf:
    # EMRFS filesystem
    fs.s3.customAWSCredentialsProvider:
com.amazonaws.auth.WebIdentityTokenCredentialsProvider
    fs.s3.impl: com.amazon.ws.emr.hadoop.fs.EmrFileSystem
    fs.AbstractFileSystem.s3.impl: org.apache.hadoop.fs.s3.EMRFSDelegate
    fs.s3.buffer.dir: /mnt/s3
    fs.s3.getObject.initialSocketTimeoutMilliseconds: "2000"

  mapreduce.fileoutputcommitter.algorithm.version.emr_internal_use_only.EmrFileSystem:
"2"
    mapreduce.fileoutputcommitter.cleanup-
failures.ignored.emr_internal_use_only.EmrFileSystem: "true"
  sparkConf:
    # Required for EMR Runtime
```



```

    spark.driver.extraClassPath: /usr/lib/hadoop-lzo/lib/*:/usr/lib/hadoop/
hadoop-aws.jar:/usr/share/aws/aws-java-sdk/*:/usr/share/aws/emr/emrfs/conf:/usr/
share/aws/emr/emrfs/lib/*:/usr/share/aws/emr/emrfs/auxlib/*:/usr/share/aws/emr/
security/conf:/usr/share/aws/emr/security/lib/*:/usr/share/aws/hmclient/lib/aws-
glue-datacatalog-spark-client.jar:/usr/share/java/Hive-JSON-Serde/hive-openx-
serde.jar:/usr/share/aws/sagemaker-spark-sdk/lib/sagemaker-spark-sdk.jar:/home/
hadoop/extrajars/*
    spark.driver.extraLibraryPath: /usr/lib/hadoop/lib/native:/usr/lib/hadoop-lzo/
lib/native:/docker/usr/lib/hadoop/lib/native:/docker/usr/lib/hadoop-lzo/lib/native
    spark.executor.extraClassPath: /usr/lib/hadoop-lzo/lib/*:/usr/lib/hadoop/
hadoop-aws.jar:/usr/share/aws/aws-java-sdk/*:/usr/share/aws/emr/emrfs/conf:/usr/
share/aws/emr/emrfs/lib/*:/usr/share/aws/emr/emrfs/auxlib/*:/usr/share/aws/emr/
security/conf:/usr/share/aws/emr/security/lib/*:/usr/share/aws/hmclient/lib/aws-
glue-datacatalog-spark-client.jar:/usr/share/java/Hive-JSON-Serde/hive-openx-
serde.jar:/usr/share/aws/sagemaker-spark-sdk/lib/sagemaker-spark-sdk.jar:/home/
hadoop/extrajars/*
    spark.executor.extraLibraryPath: /usr/lib/hadoop/lib/native:/usr/lib/hadoop-
lzo/lib/native:/docker/usr/lib/hadoop/lib/native:/docker/usr/lib/hadoop-lzo/lib/
native
    sparkVersion: "3.3.1"
    restartPolicy:
      type: Never
    driver:
      cores: 1
      coreLimit: "1200m"
      memory: "512m"
      labels:
        version: 3.3.1
      serviceAccount: my-spark-driver-sa
    executor:
      cores: 1
      instances: 1
      memory: "512m"
      labels:
        version: 3.3.1
EOF

```

If you're using the spark operator with a version 7 release, you adjust some of the configuration values:

```

cat >word-count.yaml <<EOF
apiVersion: "sparkoperator.k8s.io/v1beta2"
kind: SparkApplication

```

```

metadata:
  name: word-count
  namespace: spark-operator
spec:
  type: Java
  mode: cluster
  image: "895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/emr-7.7.0:latest"
  imagePullPolicy: Always
  mainClass: org.apache.spark.examples.JavaWordCount
  mainApplicationFile: local:///usr/lib/spark/examples/jars/spark-examples.jar
  arguments:
    - s3://my-pod-bucket/poem.txt
  hadoopConf:
    # EMRFS filesystem
    fs.s3.customAWSCredentialsProvider:
com.amazonaws.auth.WebIdentityTokenCredentialsProvider
    fs.s3.impl: com.amazon.ws.emr.hadoop.fs.EmrFileSystem
    fs.AbstractFileSystem.s3.impl: org.apache.hadoop.fs.s3.EMRFSDelegate
    fs.s3.buffer.dir: /mnt/s3
    fs.s3.getObject.initialSocketTimeoutMilliseconds: "2000"

  mapreduce.fileoutputcommitter.algorithm.version.emr_internal_use_only.EmrFileSystem:
"2"
    mapreduce.fileoutputcommitter.cleanup-
failures.ignored.emr_internal_use_only.EmrFileSystem: "true"
  sparkConf:
    # Required for EMR Runtime
    spark.driver.extraClassPath: /usr/lib/hadoop-lzo/lib/*:/usr/lib/hadoop/hadoop-
aws.jar:/usr/share/aws/aws-java-sdk/*:/usr/share/aws/aws-java-sdk-v2/*:/usr/share/
aws/emr/emrfs/conf:/usr/share/aws/emr/emrfs/lib/*:/usr/share/aws/emr/emrfs/auxlib/
*/usr/share/aws/emr/security/conf:/usr/share/aws/emr/security/lib/*:/usr/share/
aws/hmclient/lib/aws-glue-datacatalog-spark-client.jar:/usr/share/java/Hive-JSON-
Serde/hive-openx-serde.jar:/usr/share/aws/sagemaker-spark-sdk/lib/sagemaker-spark-
sdk.jar:/home/hadoop/extrajars/*
    spark.driver.extraLibraryPath: /usr/lib/hadoop/lib/native:/usr/lib/hadoop-lzo/
lib/native:/docker/usr/lib/hadoop/lib/native:/docker/usr/lib/hadoop-lzo/lib/native
    spark.executor.extraClassPath: /usr/lib/hadoop-lzo/lib/*:/usr/lib/hadoop/
hadoop-aws.jar:/usr/share/aws/aws-java-sdk/*:/usr/share/aws/aws-java-sdk-v2/*:/usr/
share/aws/emr/emrfs/conf:/usr/share/aws/emr/emrfs/lib/*:/usr/share/aws/emr/emrfs/
auxlib/*:/usr/share/aws/emr/security/conf:/usr/share/aws/emr/security/lib/*:/usr/
share/aws/hmclient/lib/aws-glue-datacatalog-spark-client.jar:/usr/share/java/Hive-
JSON-Serde/hive-openx-serde.jar:/usr/share/aws/sagemaker-spark-sdk/lib/sagemaker-
spark-sdk.jar:/home/hadoop/extrajars/*

```

```

    spark.executor.extraLibraryPath: /usr/lib/hadoop/lib/native:/usr/lib/hadoop-
lzo/lib/native:/docker/usr/lib/hadoop/lib/native:/docker/usr/lib/hadoop-lzo/lib/
native
    sparkVersion: "3.3.1"
    restartPolicy:
      type: Never
    driver:
      cores: 1
      coreLimit: "1200m"
      memory: "512m"
      labels:
        version: 3.3.1
      serviceAccount: my-spark-driver-sa
    executor:
      cores: 1
      instances: 1
      memory: "512m"
      labels:
        version: 3.3.1
EOF

```

2. Submit the Spark application.

```
kubectl apply -f word-count.yaml
```

The `kubectl` command should return confirmation that you successfully created a `SparkApplication` object called `word-count`.

```
sparkapplication.sparkoperator.k8s.io/word-count configured
```

3. To check events for the `SparkApplication` object, run the following command:

```
kubectl describe sparkapplication word-count -n spark-operator
```

The `kubectl` command should return the description of the `SparkApplication` with the events:

```

Events:
  Type          Reason          Age          From
  Message

```

```

-----
Normal    SparkApplicationSpecUpdateProcessed    3m2s (x2 over 17h)    spark-
operator  Successfully processed spec update for SparkApplication word-count
Warning   SparkApplicationPendingRerun           3m2s (x2 over 17h)    spark-
operator  SparkApplication word-count is pending rerun
Normal    SparkApplicationSubmitted               2m58s (x2 over 17h)    spark-
operator  SparkApplication word-count was submitted successfully
Normal    SparkDriverRunning                    2m56s (x2 over 17h)    spark-
operator  Driver word-count-driver is running
Normal    SparkExecutorPending                  2m50s                  spark-
operator  Executor [javawordcount-fdd1698807392c66-exec-1] is pending
Normal    SparkExecutorRunning                  2m48s                  spark-
operator  Executor [javawordcount-fdd1698807392c66-exec-1] is running
Normal    SparkDriverCompleted                  2m31s (x2 over 17h)    spark-
operator  Driver word-count-driver completed
Normal    SparkApplicationCompleted              2m31s (x2 over 17h)    spark-
operator  SparkApplication word-count completed
Normal    SparkExecutorCompleted                 2m31s (x2 over 2m31s)  spark-
operator  Executor [javawordcount-fdd1698807392c66-exec-1] completed

```

The application is now counting the words in your S3 file. To find the count of words, refer to the log files for your driver:

```
kubectl logs pod/word-count-driver -n spark-operator
```

The kubectl command should return the contents of the log file with the results of your word-count application.

```
INFO DAGScheduler: Job 0 finished: collect at JavaWordCount.java:53, took 5.146519 s
Software: 1
```

For more information on how to submit applications to Spark through the Spark operator, see [Using a SparkApplication](#) in the *Kubernetes Operator for Apache Spark (spark-on-k8s-operator)* documentation on GitHub.

Running Spark jobs with spark-submit

Amazon EMR releases 6.10.0 and higher support `spark-submit` as a command-line tool that you can use to submit and execute Spark applications to an Amazon EMR on EKS cluster.

Note

Amazon EMR calculates pricing on Amazon EKS based on vCPU and memory consumption. This calculation applies to driver and executor pods. This calculation starts from when you download your Amazon EMR application image until the Amazon EKS pod terminates and is rounded to the nearest second. Pricing applies whenever pods run using an image derived from the Amazon EMR runtime, regardless of whether the image is stored in the public Amazon EMR Amazon ECR registry or your own private registry, and regardless of how the job is orchestrated (for example, the StartJobRun API, the Spark operator, or spark-submit).

Topics

- [Setting up spark-submit for Amazon EMR on EKS](#)
- [Getting started with spark-submit for Amazon EMR on EKS](#)
- [Verify Spark driver service account security requirements for spark-submit](#)

Setting up spark-submit for Amazon EMR on EKS

Complete the following tasks to get set up before you can run an application with spark-submit on Amazon EMR on EKS. If you've already signed up for Amazon Web Services (AWS) and have used Amazon EKS, you are almost ready to use Amazon EMR on EKS. If you've already completed any of the prerequisites, you can skip those and move on to the next one.

- [Install or update to the latest version of the AWS CLI](#) – If you've already installed the AWS CLI, confirm that you have the latest version.
- [Set up kubectl and eksctl](#) – eksctl is a command line tool that you use to communicate with Amazon EKS.
- [Get started with Amazon EKS – eksctl](#) – Follow the steps to create a new Kubernetes cluster with nodes in Amazon EKS.
- [Select an Amazon EMR base image URI \(release 6.10.0 or higher\)](#) – the spark-submit command is supported with Amazon EMR releases 6.10.0 and higher.
- Confirm that the driver service account has appropriate permissions to create and watch executor pods. For more information, see [Verify Spark driver service account security requirements for spark-submit](#).

- Set up your local [AWS credentials profile](#).
- From the Amazon EKS console, choose your EKS cluster, then find the EKS cluster endpoint, located under **Overview**, **Details**, then **API server endpoint**.

Getting started with spark-submit for Amazon EMR on EKS

Amazon EMR 6.10.0 and higher supports spark-submit for running Spark applications on an Amazon EKS cluster. The section that follows shows you how to submit a command for a Spark application.

Run a Spark application

To run the Spark application, follow these steps:

1. Before you can run a Spark application with the spark-submit command, complete the steps in [Setting up spark-submit for Amazon EMR on EKS](#).
2. Run a container with an Amazon EMR on EKS base image. See [How to select a base image URI](#) for more information.

```
kubectl run -it containerName --image=EMRonEKSIImage --command -n namespace /bin/  
bash
```

3. Set the values for the following environment variables:

```
export SPARK_HOME=spark-home  
export MASTER_URL=k8s://Amazon EKS-cluster-endpoint
```

4. Now, submit the Spark application with the following command:

```
$SPARK_HOME/bin/spark-submit \  
  --class org.apache.spark.examples.SparkPi \  
  --master $MASTER_URL \  
  --conf spark.kubernetes.container.image=895885662937.dkr.ecr.us-  
west-2.amazonaws.com/spark/emr-6.10.0:latest \  
  --conf spark.kubernetes.authenticate.driver.serviceAccountName=spark \  
  --deploy-mode cluster \  
  --conf spark.kubernetes.namespace=spark-operator \  
  local:///usr/lib/spark/examples/jars/spark-examples.jar 20
```

For more information about submitting applications to Spark, see [Submitting applications](#) in the Apache Spark documentation.

⚠ Important

`spark-submit` only supports cluster mode as the submission mechanism.

Verify Spark driver service account security requirements for `spark-submit`

The Spark driver pod uses a Kubernetes service account to access the Kubernetes API server to create and watch executor pods. Driver service account must have appropriate permissions to list, create, edit, patch and delete pods in your cluster. You can verify that you can list these resources by running the following command:

```
kubectl auth can-i list/create/edit/delete/patch pods
```

Verify that you have the necessary permissions by running each command.

```
kubectl auth can-i list pods
kubectl auth can-i create pods
kubectl auth can-i edit pods
kubectl auth can-i delete pods
kubectl auth can-i patch pods
```

The following rules apply to this service role:

```
rules:
- apiGroups:
  - ""
  resources:
  - pods
  verbs:
  - "*"
- apiGroups:
  - ""
  resources:
  - services
  verbs:
```

```

- "*"
- apiGroups:
  - ""
  resources:
  - configmaps
  verbs:
  - "*"
- apiGroups:
  - ""
  resources:
  - persistentvolumeclaims
  verbs:
  - "*"

```

Setting up IAM roles for service accounts (IRSA) for spark-submit

The following sections explain how to set up IAM roles for service accounts (IRSA) to authenticate and authorize Kubernetes service accounts so you can run Spark applications stored in Amazon S3.

Prerequisites

Before trying any of the examples in this documentation, make sure that you have completed the following prerequisites:

- [Finished setting up spark-submit](#)
- [Created an S3 bucket](#) and [uploaded](#) the spark application jar

Configuring a Kubernetes service account to assume an IAM role

The following steps cover how to configure a Kubernetes service account to assume an AWS Identity and Access Management (IAM) role. After you configure the pods to use the service account, they can then access any AWS service that the role has permissions to access.

1. Create a policy file to allow read-only access to the Amazon S3 object you [uploaded](#):

```

cat >my-policy.json <<EOF
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",

```



```

        "Action": [
            "s3:GetObject",
            "s3:ListBucket"
        ],
        "Resource": [
            "arn:aws:s3:::<my-spark-jar-bucket>",
            "arn:aws:s3:::<my-spark-jar-bucket>/*"
        ]
    }
]
}
EOF

```

2. Create the IAM policy.

```
aws iam create-policy --policy-name my-policy --policy-document file://my-policy.json
```

3. Create an IAM role and associate it with a Kubernetes service account for the Spark driver

```
eksctl create iamserviceaccount --name my-spark-driver-sa --namespace spark-operator \
--cluster my-cluster --role-name "my-role" \
--attach-policy-arn arn:aws:iam::111122223333:policy/my-policy --approve
```

4. Create a YAML file with the required [permissions](#) for the Spark driver service account:

```
cat >spark-rbac.yaml <<EOF
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  namespace: default
  name: emr-containers-role-spark
rules:
- apiGroups:
  - ""
  resources:
  - pods
  verbs:
  - "*"
- apiGroups:
  - ""
  resources:

```

```
- services
verbs:
- "*"
- apiGroups:
- ""
resources:
- configmaps
verbs:
- "*"
- apiGroups:
- ""
resources:
- persistentvolumeclaims
verbs:
- "*"
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: spark-role-binding
  namespace: default
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: Role
  name: emr-containers-role-spark
subjects:
- kind: ServiceAccount
  name: emr-containers-sa-spark
  namespace: default
EOF
```

5. Apply the cluster role binding configurations.

```
kubectl apply -f spark-rbac.yaml
```

6. The `kubectl` command should return confirmation of the created account.

```
serviceaccount/emr-containers-sa-spark created
clusterrolebinding.rbac.authorization.k8s.io/emr-containers-role-spark configured
```

Running the Spark application

Amazon EMR 6.10.0 and higher supports spark-submit for running Spark applications on an Amazon EKS cluster. To run the Spark application, follow these steps:

1. Make sure that you have completed the steps in [Setting up spark-submit for Amazon EMR on EKS](#).
2. Set the values for the following environment variables:

```
export SPARK_HOME=spark-home
export MASTER_URL=k8s://Amazon EKS-cluster-endpoint
```

3. Now, submit the Spark application with the following command:

```
$SPARK_HOME/bin/spark-submit \
  --class org.apache.spark.examples.SparkPi \
  --master $MASTER_URL \
  --conf spark.kubernetes.container.image=895885662937.dkr.ecr.us-
west-2.amazonaws.com/spark/emr-6.15.0:latest \
  --conf spark.kubernetes.authenticate.driver.serviceAccountName=emr-containers-sa-
spark \
  --deploy-mode cluster \
  --conf spark.kubernetes.namespace=default \
  --conf "spark.driver.extraClassPath=/usr/lib/hadoop-lzo/lib/*:/usr/lib/hadoop/
hadoop-aws.jar:/usr/share/aws/aws-java-sdk/*:/usr/share/aws/emr/emrfs/conf:/usr/
share/aws/emr/emrfs/lib/*:/usr/share/aws/emr/emrfs/auxlib/*:/usr/share/aws/emr/
security/conf:/usr/share/aws/emr/security/lib/*:/usr/share/aws/hmclient/lib/aws-
glue-datacatalog-spark-client.jar:/usr/share/java/Hive-JSON-Serde/hive-openx-
serde.jar:/usr/share/aws/sagemaker-spark-sdk/lib/sagemaker-spark-sdk.jar:/home/
hadoop/extrajars/*" \
  --conf "spark.driver.extraLibraryPath=/usr/lib/hadoop/lib/native:/usr/lib/hadoop-
lzo/lib/native:/docker/usr/lib/hadoop/lib/native:/docker/usr/lib/hadoop-lzo/lib/
native" \
  --conf "spark.executor.extraClassPath=/usr/lib/hadoop-lzo/lib/*:/usr/lib/hadoop/
hadoop-aws.jar:/usr/share/aws/aws-java-sdk/*:/usr/share/aws/emr/emrfs/conf:/usr/
share/aws/emr/emrfs/lib/*:/usr/share/aws/emr/emrfs/auxlib/*:/usr/share/aws/emr/
security/conf:/usr/share/aws/emr/security/lib/*:/usr/share/aws/hmclient/lib/aws-
glue-datacatalog-spark-client.jar:/usr/share/java/Hive-JSON-Serde/hive-openx-
serde.jar:/usr/share/aws/sagemaker-spark-sdk/lib/sagemaker-spark-sdk.jar:/home/
hadoop/extrajars/*" \
  --conf "spark.executor.extraLibraryPath=/usr/lib/hadoop/lib/native:/usr/lib/
hadoop-lzo/lib/native:/docker/usr/lib/hadoop/lib/native:/docker/usr/lib/hadoop-lzo/
lib/native" \
```

```
--conf
spark.hadoop.fs.s3.customAWSCredentialsProvider=com.amazonaws.auth.WebIdentityTokenCredent
\
--conf spark.hadoop.fs.s3.impl=com.amazon.ws.emr.hadoop.fs.EmrFileSystem \
--conf
spark.hadoop.fs.AbstractFileSystem.s3.impl=org.apache.hadoop.fs.s3.EMRFSDelegate \
--conf spark.hadoop.fs.s3.buffer.dir=/mnt/s3 \
--conf spark.hadoop.fs.s3.getObject.initialSocketTimeoutMilliseconds="2000" \
--conf
spark.hadoop.mapreduce.fileoutputcommitter.algorithm.version.emr_internal_use_only.EmrFile
\
--conf spark.hadoop.mapreduce.fileoutputcommitter.cleanup-
failures.ignored.emr_internal_use_only.EmrFileSystem="true" \
s3://my-pod-bucket/spark-examples.jar 20
```

4. After the spark driver finishes the Spark job, you should see a log line at the end of the submission indicating that the Spark job has finished.

```
23/11/24 17:02:14 INFO LoggingPodStatusWatcherImpl: Application
org.apache.spark.examples.SparkPi with submission ID default:org-apache-spark-
examples-sparkpi-4980808c03ff3115-driver finished
23/11/24 17:02:14 INFO ShutdownHookManager: Shutdown hook called
```

Cleanup

When you're done running your applications, you can perform cleanup with the following command.

```
kubectl delete -f spark-rbac.yaml
```

Using Apache Livy with Amazon EMR on EKS

With Amazon EMR releases 7.1.0 and higher, you can use Apache Livy to submit jobs on Amazon EMR on EKS. Using Apache Livy, you can set up your own Apache Livy REST endpoint and use it to deploy and manage Spark applications on your Amazon EKS clusters. After you install Livy in your Amazon EKS cluster, you can use the Livy endpoint to submit Spark applications to your Livy server. The server manages the lifecycle of the Spark applications.

Note

Amazon EMR calculates pricing on Amazon EKS based on vCPU and memory consumption. This calculation applies to driver and executor pods. This calculation starts from when you download your Amazon EMR application image until the Amazon EKS pod terminates and is rounded to the nearest second. Pricing applies whenever pods run using an image derived from the Amazon EMR runtime, regardless of whether the image is stored in the public Amazon EMR Amazon ECR registry or your own private registry, and regardless of how the job is orchestrated (for example, the StartJobRun API, the Spark operator, or spark-submit).

Topics

- [Setting up Apache Livy for Amazon EMR on EKS](#)
- [Getting started with Apache Livy on Amazon EMR on EKS](#)
- [Running a Spark application with Apache Livy for Amazon EMR on EKS](#)
- [Uninstalling Apache Livy with Amazon EMR on EKS](#)
- [Security for Apache Livy with Amazon EMR on EKS](#)
- [Installation properties for Apache Livy on Amazon EMR on EKS releases](#)
- [Troubleshoot common environment-variable format errors](#)

Setting up Apache Livy for Amazon EMR on EKS

Before you can install Apache Livy on your Amazon EKS cluster, you must install and configure a set of prerequisite tools. These include the AWS CLI, which is a foundational command-line tool for working with AWS resources, command-line tools for working with Amazon EKS, and a controller that's used in this use case to make your cluster application available to the internet and to route network traffic.

- [Install or update to the latest version of the AWS CLI](#) – If you've already installed the AWS CLI, confirm that you have the latest version.
- [Set up kubectl and eksctl](#) – eksctl is a command line tool that you use to communicate with Amazon EKS.
- [Install Helm](#) – The Helm package manager for Kubernetes helps you install and manage applications on your Kubernetes cluster.

- [Get started with Amazon EKS – eksctl](#) – Follow the steps to create a new Kubernetes cluster with nodes in Amazon EKS.
- [Select an Amazon EMR release label](#) – the Apache Livy is supported with Amazon EMR releases 7.1.0 and higher.
- [Install the ALB controller](#) – the ALB controller manages AWS Elastic Load Balancing for Kubernetes clusters. It creates an AWS Network Load Balancer (NLB) when you create a Kubernetes Ingress while setting up Apache Livy.

Getting started with Apache Livy on Amazon EMR on EKS

Complete the following steps to install Apache Livy. They include configuring the package manager, creating a namespace for running Spark workloads, installing Livy, setting up load balancing, and verification steps. You have to complete these steps in order to run a batch job with Spark.

1. If you haven't already, set up [Apache Livy for Amazon EMR on EKS](#).
2. Authenticate your Helm client to the Amazon ECR registry. You can find the corresponding ECR-registry-account value for your AWS Region from [Amazon ECR registry accounts by Region](#).

```
aws ecr get-login-password --region <AWS_REGION> | helm registry login \
--username AWS \
--password-stdin <ECR-registry-account>.dkr.ecr.<region-id>.amazonaws.com
```

3. Setting up Livy creates a service account for the Livy server and another account for the Spark application. To set up IRSA for the service accounts, see [Setting up access permissions with IAM roles for service accounts \(IRSA\)](#).
4. Create a namespace to run your Spark workloads.

```
kubectl create ns <spark-ns>
```

5. Use the following command to install Livy.

This Livy endpoint is only internally available to the VPC in the EKS cluster. To enable access beyond the VPC, set `--set loadbalancer.internal=false` in your Helm installation command.

Note

By default, SSL is not enabled within this Livy endpoint and the endpoint is only visible inside the VPC of the EKS cluster. If you set `loadbalancer.internal=false` and `ssl.enabled=false`, you are exposing an insecure endpoint to outside of your VPC. To set up a secure Livy endpoint, see [Configuring a secure Apache Livy endpoint with TLS/SSL](#).

```
helm install livy-demo \
  oci://895885662937.dkr.ecr.region-id.amazonaws.com/livy \
  --version 7.13.0 \
  --namespace livy-ns \
  --set image=ECR-registry-account.dkr.ecr.region-id.amazonaws.com/livy/
emr-7.13.0:latest \
  --set sparkNamespace=<spark-ns> \
  --create-namespace
```

You should see the following output.

```
NAME: livy-demo
LAST DEPLOYED: Mon Mar 18 09:23:23 2024
NAMESPACE: livy-ns
STATUS: deployed
REVISION: 1
TEST SUITE: None
NOTES:
The Livy server has been installed.
Check installation status:
1. Check Livy Server pod is running
   kubectl --namespace livy-ns get pods -l "app.kubernetes.io/instance=livy-demo"
2. Verify created NLB is in Active state and it's target groups are healthy (if
   loadbalancer.enabled is true)

Access LIVY APIs:
  # Ensure your NLB is active and healthy
  # Get the Livy endpoint using command:
  LIVY_ENDPOINT=$(kubectl get svc -n livy-ns -l app.kubernetes.io/
instance=livy-demo,emr-containers.amazonaws.com/type=loadbalancer -o
```

```
jsonpath='{.items[0].status.loadBalancer.ingress[0].hostname}' | awk '{printf
"%s:8998\n", $0}'
# Access Livy APIs using http://$LIVY_ENDPOINT or https://$LIVY_ENDPOINT (if
SSL is enabled)
# Note: While uninstalling Livy, makes sure the ingress and NLB are deleted
after running the helm command to avoid dangling resources
```

The default service account names for the Livy server and the Spark session are `emr-containers-sa-livy` and `emr-containers-sa-spark-livy`. To use custom names, use the `serviceAccounts.name` and `sparkServiceAccount.name` parameters.

```
--set serviceAccounts.name=my-service-account-for-livy
--set sparkServiceAccount.name=my-service-account-for-spark
```

6. Verify that you installed the Helm chart.

```
helm list -n livy-ns -o yaml
```

The `helm list` command should return information about your new Helm chart.

```
app_version: 0.7.1-incubating
chart: livy-emr-7.13.0
name: livy-demo
namespace: livy-ns
revision: "1"
status: deployed
updated: 2024-02-08 22:39:53.539243 -0800 PST
```

7. Verify that the Network Load Balancer is active.

```
LIVY_NAMESPACE=<livy-ns>
LIVY_APP_NAME=<livy-app-name>
AWS_REGION=<AWS_REGION>

# Get the NLB Endpoint URL
NLB_ENDPOINT=$(kubectl --namespace $LIVY_NAMESPACE get svc -l "app.kubernetes.io/
instance=$LIVY_APP_NAME,emr-containers.amazonaws.com/type=loadbalancer" -o
jsonpath='{.items[0].status.loadBalancer.ingress[0].hostname}')

# Get all the load balancers in the account's region
ELB_LIST=$(aws elbv2 describe-load-balancers --region $AWS_REGION)
```



```
# Get the status of the NLB that matching the endpoint from the Kubernetes service
NLB_STATUS=$(echo $ELB_LIST | grep -A 8 "\"DNSName\": \"$NLB_ENDPOINT\"" | awk '/
Code/{print $2}/' | tr -d '",'')
echo $NLB_STATUS
```

8. Now verify that the target group in the Network Load Balancer is healthy.

```
LIVY_NAMESPACE=<livy-ns>
LIVY_APP_NAME=<livy-app-name>
AWS_REGION=<AWS_REGION>

# Get the NLB endpoint
NLB_ENDPOINT=$(kubectl --namespace $LIVY_NAMESPACE get svc -l "app.kubernetes.io/
instance=$LIVY_APP_NAME,emr-containers.amazonaws.com/type=loadbalancer" -o
jsonpath='{.items[0].status.loadBalancer.ingress[0].hostname}')

# Get all the load balancers in the account's region
ELB_LIST=$(aws elbv2 describe-load-balancers --region $AWS_REGION)

# Get the NLB ARN from the NLB endpoint
NLB_ARN=$(echo $ELB_LIST | grep -B 1 "\"DNSName\": \"$NLB_ENDPOINT\"" | awk
'/"LoadBalancerArn":/,/,/' | awk '/:/{print $2}' | tr -d \,)

# Get the target group from the NLB. Livy setup only deploys 1 target group
TARGET_GROUP_ARN=$(aws elbv2 describe-target-groups --load-balancer-arn $NLB_ARN
--region $AWS_REGION | awk '/"TargetGroupArn":/,/,/' | awk '/:/{print $2}' | tr -d
\,,)

# Get health of target group
aws elbv2 describe-target-health --target-group-arn $TARGET_GROUP_ARN
```

The following is sample output that shows the status of the target group:

```
{
  "TargetHealthDescriptions": [
    {
      "Target": {
        "Id": "<target IP>",
        "Port": 8998,
        "AvailabilityZone": "us-west-2d"
      },
      "HealthCheckPort": "8998",
```

```

        "TargetHealth": {
            "State": "healthy"
        }
    }
]
}

```

Once the status of your NLB becomes active and your target group is healthy, you can continue. It might take a few minutes.

9. Retrieve the Livy endpoint from the Helm installation. Whether or not your Livy endpoint is secure depends on whether you enabled SSL.

```

LIVY_NAMESPACE=<livy-ns>
LIVY_APP_NAME=<livy-app-name>
LIVY_ENDPOINT=$(kubectl get svc -n livy-ns -l app.kubernetes.io/
instance=<livy-app-name>,emr-containers.amazonaws.com/type=loadbalancer -o
jsonpath='{.items[0].status.loadBalancer.ingress[0].hostname}' | awk '{printf
"%s:8998\n", $0}')
echo "$LIVY_ENDPOINT"

```

10. Retrieve the Spark service account from the Helm installation

```

SPARK_NAMESPACE=spark-ns
LIVY_APP_NAME=<livy-app-name>
SPARK_SERVICE_ACCOUNT=$(kubectl --namespace $SPARK_NAMESPACE
get sa -l "app.kubernetes.io/instance=$LIVY_APP_NAME" -o
jsonpath='{.items[0].metadata.name}')
echo "$SPARK_SERVICE_ACCOUNT"

```

You should see something similar to the following output:

```
emr-containers-sa-spark-livy
```

11. If you set `internalALB=true` to enable access from outside of your VPC, create an Amazon EC2 instance and make sure the Network Load Balancer allows network traffic coming from the EC2 instance. You must do so for the instance to have access to your Livy endpoint. For more information about securely exposing your endpoint outside of your VPC, see [Setting up with a secure Apache Livy endpoint with TLS/SSL](#).

12. Installing Livy creates the service account `emr-containers-sa-spark` to run Spark applications. If your Spark application uses any AWS resources like S3 or calls AWS API or CLI operations, you must link an IAM role with the necessary permissions to your spark service account. For more information, see [Setting up access permissions with IAM roles for service accounts \(IRSA\)](#).

Apache Livy supports additional configurations that you can use while installing Livy. For more information, see Installation properties for Apache Livy on Amazon EMR on EKS releases.

Running a Spark application with Apache Livy for Amazon EMR on EKS

Before you can run a Spark application with Apache Livy, make sure that you have completed the steps in [Setting up Apache Livy for Amazon EMR on EKS](#) and [Getting started with Apache Livy for Amazon EMR on EKS](#).

You can use Apache Livy to run two types of applications:

- Batch sessions – a type of Livy workload to submit Spark batch jobs.
- Interactive sessions – a type of Livy workload that provides a programmatic and visual interface to run Spark queries.

Note

Driver and executor pods from different sessions can communicate with each other. Namespaces don't guarantee any security between pods. Kubernetes doesn't allow selective permissions on a subset of pods inside a given namespace.

Running batch sessions

To submit a batch job, use the following command.

```
curl -s -k -H 'Content-Type: application/json' -X POST \
  -d '{
    "name": "my-session",
    "file": "entryPoint_location (S3 or local)",
    "args": ["argument1", "argument2", ...],
    "conf": {
      "spark.kubernetes.namespace": "<spark-namespace>",
```

```

        "spark.kubernetes.container.image": "public.ecr.aws/emr-on-eks/spark/
emr-7.13.0:latest",
        "spark.kubernetes.authenticate.driver.serviceAccountName": "<spark-
service-account>"
    }
} ' <livy-endpoint>/batches

```

To monitor your batch job, use the following command.

```

curl -s -k -H 'Content-Type: application/json' -X GET <livy-endpoint>/batches/my-
session

```

Running interactive sessions

To run interactive sessions with Apache Livy, see the following steps.

1. Make sure you have access to either a self-hosted or a managed Jupyter notebook, such as a SageMaker AI Jupyter notebook. Your jupyter notebook must have [sparkmagic](#) installed.
2. Create a bucket for Spark configuration `spark.kubernetes.file.upload.path`. Make sure the Spark service account has read and write access to the bucket. For more details on how to configure your spark service account, see [Setting up access permissions with IAM roles for service accounts \(IRSA\)](#)
3. Load sparkmagic in the Jupyter notebook with the command `%load_ext sparkmagic.magics`.
4. Run the command `%manage_spark` to set up your Livy endpoint with the Jupyter notebook. Choose the **Add Endpoints** tab, choose the configured auth type, add the Livy endpoint to the notebook, and then choose **Add endpoint**.
5. Run `%manage_spark` again to create the Spark context and then go to the **Create session**. Choose the Livy endpoint, specify a unique session name choose a language, and then add the following properties.

```

{
  "conf": {
    "spark.kubernetes.namespace": "livy-namespace",
    "spark.kubernetes.container.image": "public.ecr.aws/emr-on-eks/spark/
emr-7.13.0:latest",
    "spark.kubernetes.authenticate.driver.serviceAccountName": "<spark-service-
account>",
    "spark.kubernetes.file.upload.path": "<URI_TO_S3_LOCATION>"
  }
}

```

```
}  
}
```

6. Submit the application and wait for it to create the Spark context.
7. To monitor the status of the interactive session, run the following command.

```
curl -s -k -H 'Content-Type: application/json' -X GET livy-endpoint/sessions/my-interactive-session
```

Monitoring Spark applications

To monitor the progress of your Spark applications with the Livy UI, use the link `http://<livy-endpoint>/ui`.

Uninstalling Apache Livy with Amazon EMR on EKS

Follow these steps to uninstall Apache Livy.

1. Delete the Livy setup using the names of your namespace and application name. In this example, the application name is `livy-demo` and the namespace is `livy-ns`.

```
helm uninstall livy-demo -n livy-ns
```

2. When uninstalling, Amazon EMR on EKS deletes the Kubernetes service in Livy, the AWS load balancers, and the target groups that you created during installation. Deleting resources can take a few minutes. Make sure that the resources are deleted before installing Livy on the namespace again.
3. Delete the Spark namespace.

```
kubectl delete namespace spark-ns
```

Security for Apache Livy with Amazon EMR on EKS

See the following topics to learn more about configuring security for Apache Livy with Amazon EMR on EKS. These options include using transport-layer security, role-based access control, which is access based on a person's role within an organization, and using IAM roles, which provide access to resources, based on granted permissions.

Topics

- [Setting up a secure Apache Livy endpoint with TLS/SSL](#)
- [Setting up the Apache Livy and Spark application permissions with role-based access control \(RBAC\)](#)
- [Setting up access permissions with IAM roles for service accounts \(IRSA\)](#)

Setting up a secure Apache Livy endpoint with TLS/SSL

See the following sections to learn more about setting up Apache Livy for Amazon EMR on EKS with end-to-end TLS and SSL encryption.

Setting up TLS and SSL encryption

To set up SSL encryption on your Apache Livy endpoint, follow these steps.

- [Install the Secrets Store CSI Driver and AWS Secrets and Configuration Provider \(ASCP\)](#) – the Secrets Store CSI Driver and ASCP securely store Livy's JKS certificates and passwords that the Livy server pod needs to enable SSL. You can also install just the Secrets Store CSI Driver and use any other supported secrets provider.
- [Create an ACM certificate](#) – this certificate is required to secure the connection between the client and the ALB endpoint.
- Set up a JKS certificate, key password, and keystore password for AWS Secrets Manager – required to secure the connection between the ALB endpoint and the Livy server.
- Add permissions to the Livy service account to retrieve secrets from AWS Secrets Manager – the Livy server needs these permissions to retrieve secrets from ASCP and add the Livy configurations to secure the Livy server. To add IAM permissions to a service account, see [Setting up access permissions with IAM roles for service accounts \(IRSA\)](#).

Setting up a JKS certificate with a key and a keystore password for AWS Secrets Manager

Follow these steps to set up a JKS certificate with a key and a keystore password.

1. Generate a keystore file for the Livy server.

```
keytool -genkey -alias <host> -keyalg RSA -keysize 2048 -dname  
CN=<host>,OU=hw,O=hw,L=<your_location>,ST=<state>,C=<country> -
```

```
keypass <keyPassword> -keystore <keystore_file> -storepass <storePassword> --
validity 3650
```

2. Create a certificate.

```
keytool -export -alias <host> -keystore mykeystore.jks -rfc -
file mycertificate.cert -storepass <storePassword>
```

3. Create a truststore file.

```
keytool -import -noprompt -alias <host>-file <cert_file> -
keystore <truststore_file> -storepass <truststorePassword>
```

4. Save the JKS certificate in AWS Secrets Manager. Replace `livy-jks-secret` with your secret and `fileb://mykeystore.jks` with the path to your keystore JKS certificate.

```
aws secretsmanager create-secret \
--name livy-jks-secret \
--description "My Livy keystore JKS secret" \
--secret-binary fileb://mykeystore.jks
```

5. Save the keystore and key password in Secrets Manager. Make sure to use your own parameters.

```
aws secretsmanager create-secret \
--name livy-jks-secret \
--description "My Livy key and keystore password secret" \
--secret-string "{\"keyPassword\": \"<test-key-password>\", \"keyStorePassword\": \"<test-key-store-password>\"}"
```

6. Create a Livy server namespace with the following command.

```
kubectl create ns <livy-ns>
```

7. Create the `ServiceProviderClass` object for the Livy server that has the JKS certificate and the passwords.

```
cat >livy-secret-provider-class.yaml << EOF
apiVersion: secrets-store.csi.x-k8s.io/v1
kind: SecretProviderClass
metadata:
  name: aws-secrets
```

```
spec:
  provider: aws
  parameters:
    objects: |
      - objectName: "livy-jks-secret"
        objectType: "secretsmanager"
      - objectName: "livy-passwords"
        objectType: "secretsmanager"

EOF
kubectl apply -f livy-secret-provider-class.yaml -n <livy-ns>
```

Getting started with SSL-enabled Apache Livy

After enabling SSL on your Livy server, you must set up the serviceAccount to have access to the keyStore and keyPasswords secrets on AWS Secrets Manager.

1. Create the Livy server namespace.

```
kubectl create namespace <livy-ns>
```

2. Set up the Livy service account to have access to the secrets in Secrets Manager. For more information about setting up IRSA, see [Setting up IRSA while installing Apache Livy](#).

```
aws ecr get-login-password --region region-id | helm registry login \
--username AWS \
--password-stdin ECR-registry-account.dkr.ecr.region-id.amazonaws.com
```

3. Install Livy. For the Helm chart --version parameter, use your Amazon EMR release label, such as 7.1.0. You must also replace the Amazon ECR registry account ID and Region ID with your own IDs. You can find the corresponding ECR-registry-account value for your AWS Region from [Amazon ECR registry accounts by Region](#).

```
helm install <livy-app-name> \
oci://895885662937.dkr.ecr.region-id.amazonaws.com/livy \
--version 7.13.0 \
--namespace livy-namespace-name \
--set image=<ECR-registry-account.dkr.ecr>.<region>.amazonaws.com/livy/
emr-7.13.0:latest \
--set sparkNamespace=spark-namespace \
--set ssl.enabled=true
```



```
--set ssl.CertificateArn=livy-acm-certificate-arn
--set ssl.secretProviderClassName=aws-secrets
--set ssl.keyStoreObjectName=livy-jks-secret
--set ssl.keyPasswordsObjectName=livy-passwords
--create-namespace
```

4. Continue from step 5 of the [Installing Apache Livy on Amazon EMR on EKS](#).

Setting up the Apache Livy and Spark application permissions with role-based access control (RBAC)

To deploy Livy, Amazon EMR on EKS creates a server service account and role and a Spark service account and role. These roles must have the necessary RBAC permissions to finish setup and run Spark applications.

RBAC permissions for the server service account and role

Amazon EMR on EKS creates the Livy server service account and role to manage Livy sessions for Spark jobs and routing traffic to and from the ingress and other resources.

The default name for this service account is `emr-containers-sa-livy`. It must have the following permissions.

```
rules:
- apiGroups:
  - ""
  resources:
  - "namespaces"
  verbs:
  - "get"
- apiGroups:
  - ""
  resources:
  - "serviceaccounts"
  - "services"
  - "configmaps"
  - "events"
  - "pods"
  - "pods/log"
  verbs:
  - "get"
  - "list"
```

```

    "watch"
    "describe"
    "create"
    "edit"
    "delete"
    "deletecollection"
    "annotate"
    "patch"
    "label"
  - apiGroups:
    - ""
    resources:
    - "secrets"
    verbs:
    - "create"
      "patch"
      "delete"
      "watch"
  - apiGroups:
    - ""
    resources:
    - "persistentvolumeclaims"
    verbs:
    - "get"
      "list"
      "watch"
      "describe"
      "create"
      "edit"
      "delete"
      "annotate"
      "patch"
      "label"

```

RBAC permissions for the spark service account and role

A Spark driver pod needs a Kubernetes service account in the same namespace as the pod. This service account needs permissions to manage executor pods and any resources required by the driver pod. Unless the default service account in the namespace has the required permissions, the driver fails and exits. The following RBAC permissions are required.

```

rules:
- apiGroups:

```

```
- ""
  "batch"
  "extensions"
  "apps"
resources:
- "configmaps"
  "serviceaccounts"
  "events"
  "pods"
  "pods/exec"
  "pods/log"
  "pods/portforward"
  "secrets"
  "services"
  "persistentvolumeclaims"
  "statefulsets"
verbs:
- "create"
  "delete"
  "get"
  "list"
  "patch"
  "update"
  "watch"
  "describe"
  "edit"
  "deletecollection"
  "patch"
  "label"
```

Setting up access permissions with IAM roles for service accounts (IRSA)

By default, the Livy server and Spark application's driver and executors don't have access to AWS resources. The server service account and spark service account controls access to AWS resources for the Livy server and spark application's pods. To grant access, you need to map the service accounts with an IAM role that has the necessary AWS permissions.

You can set up IRSA mapping before you install Apache Livy, during the installation, or after you finish the installation.

Setting up IRSA while installing Apache Livy (for server service account)

Note

This mapping is supported only for the server service account.

1. Make sure that you have finished [setting up Apache Livy for Amazon EMR on EKS](#) and are in the middle of [installing Apache Livy with Amazon EMR on EKS](#).
2. Create a Kubernetes namespace for the Livy server. In this example, the name of the namespace is `livy-ns`.
3. Create an IAM policy that includes the permissions for the AWS services for which you want your pods to access. The following example creates an IAM policy of getting Amazon S3 resources for the Spark entry point.

```
cat >my-policy.json <<EOF{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3:::my-spark-entriypoint-bucket"
    }
  ]
}
EOF

aws iam create-policy --policy-name my-policy --policy-document file://my-policy.json
```

4. Use the following command to set your AWS account ID to a variable.

```
account_id=$(aws sts get-caller-identity --query "Account" --output text)
```

5. Set the OpenID Connect (OIDC) identity provider of your cluster to an environment variable.

```
oidc_provider=$(aws eks describe-cluster --name my-cluster --region $AWS_REGION --query "cluster.identity.oidc.issuer" --output text | sed -e "s/^https:\\/\\/\\//")
```

6. Set variables for the namespace and name of the service account. Be sure to use your own values.

```
export namespace=default
export service_account=my-service-account
```

7. Create a trust policy file with the following command. If you want to grant access of the role to all service accounts within a namespace, copy the following command, and replace `StringEquals` with `StringLike` and replace `$service_account` with `*`.

```
cat >trust-relationship.json <<EOF
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Federated": "arn:aws:iam::$account_id:oidc-provider/$oidc_provider"
      },
      "Action": "sts:AssumeRoleWithWebIdentity",
      "Condition": {
        "StringEquals": {
          "$oidc_provider:aud": "sts.amazonaws.com",
          "$oidc_provider:sub": "system:serviceaccount:$namespace:$service_account"
        }
      }
    }
  ]
}
EOF
```

8. Create the role.

```
aws iam create-role --role-name my-role --assume-role-policy-document file://trust-relationship.json --description "my-role-description"
```

9. Use the following Helm install command to set the `serviceAccount.executionRoleArn` to map IRSA. The following is an example of the Helm install command. You can find the corresponding `ECR-registry-account` value for your AWS Region from [Amazon ECR registry accounts by Region](#).

```
helm install livy-demo \
```

```
oci://895885662937.dkr.ecr.us-west-2.amazonaws.com/livy \
--version 7.13.0 \
--namespace livy-ns \
--set image=ECR-registry-account.dkr.ecr.region-id.amazonaws.com/livy/
emr-7.13.0:latest \
--set sparkNamespace=spark-ns \
--set serviceAccount.executionRoleArn=arn:aws:iam::123456789012:role/my-role
```

Mapping IRSA to a Spark service account

Before you map IRSA to a Spark service account, make sure that you have completed the following items:

- Make sure that you have finished [setting up Apache Livy for Amazon EMR on EKS](#) and are in the middle of [installing Apache Livy with Amazon EMR on EKS](#).
- You must have an existing IAM OpenID Connect (OIDC) provider for your cluster. To see if you already have one or how to create one, see [Create an IAM OIDC provider for your cluster](#).
- Make sure that you have installed version 0.171.0 or later of the `eksctl` CLI installed or AWS CloudShell. To install or update `eksctl`, see [Installation](#) of the `eksctl` documentation.

Follow these steps to map IRSA to your Spark service account:

1. Use the following command to get the Spark service account.

```
SPARK_NAMESPACE=<spark-ns>
LIVY_APP_NAME=<livy-app-name>
kubectl --namespace $SPARK_NAMESPACE describe sa -l "app.kubernetes.io/instance=
$LIVY_APP_NAME" | awk '/^Name:/ {print $2}'
```

2. Set your variables for the namespace and name of the service account.

```
export namespace=default
export service_account=my-service-account
```

3. Use the following command to create a trust policy file for the IAM role. The following example gives permission to all service accounts within the namespace to use the role. To do so, replace `StringEquals` with `StringLike` and replace `$service_account` with `*`.

```
cat >trust-relationship.json <<EOF
```

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Federated": "arn:aws:iam::${account_id}:oidc-provider/${oidc_provider}"
      },
      "Action": "sts:AssumeRoleWithWebIdentity",
      "Condition": {
        "StringEquals": {
          "${oidc_provider}:aud": "sts.amazonaws.com",
          "${oidc_provider}:sub": "system:serviceaccount:${namespace}:${service_account}"
        }
      }
    }
  ]
}
EOF
```

4. Create the role.

```
aws iam create-role --role-name my-role --assume-role-policy-document file://trust-relationship.json --description "my-role-description"
```

5. Map the server or spark service account with the following `eksctl` command. Make sure to use your own values.

```
eksctl create iamserviceaccount --name spark-sa \
--namespace spark-namespace --cluster livy-eks-cluster \
--attach-role-arn arn:aws:iam::0123456789012:role/my-role \
--approve --override-existing-serviceaccounts
```

Installation properties for Apache Livy on Amazon EMR on EKS releases

Apache Livy installation allows you to select a version of the Livy Helm chart. The Helm chart offers a variety of properties to customize your installation and setup experience. These properties are supported for Amazon EMR on EKS releases 7.1.0 and higher.

Topics

- [Amazon EMR 7.1.0 installation properties](#)

Amazon EMR 7.1.0 installation properties

The following table describes all of the supported Livy properties. When installing Apache Livy, you can choose the Livy Helm chart version. To set a property during the installation, use the command `--set <property>=<value>`.

Property	Description	Default
image	The Amazon EMR release URI of the Livy server. This is a required configuration.	""
sparkNamespace	Namespace to run Livy Spark sessions. For example, specify "livy". This is a required configuration.	""
nameOverride	Provide a name instead of <code>livy</code> . The name is set as a label for all Livy resources	"livy"
fullnameOverride	Provide a name to use instead of the full names of resources.	""
ssl.enabled	Enables end-to-end SSL from Livy endpoint to Livy server.	FALSE
ssl.certificateArn	If SSL is enabled, this is the ACM certificate ARN for the NLB created by the service..	""
ssl.secretProviderClassName	If SSL is enabled, this is the secret provider class name to secure NLB for the Livy server connection with SSL.	""
ssl.keyStoreObjectName	If SSL is enabled, the object name for the keystore	""

Property	Description	Default
	certificate in the secret provider class.	
ssl.keyPasswordsObjectName	If SSL is enabled, the object name for the secret that has the keystore and key password.	""
rbac.create	If true, creates RBAC resources.	FALSE
serviceAccount.create	If true, creates a Livy service account.	TRUE
serviceAccount.name	The name of the service account to use for Livy. If you don't set this property and create a service account, Amazon EMR on EKS automatically generates a name using the <code>fullname</code> override property.	"emr-containers-sa-livy"
serviceAccount.executionRoleArn	The execution role ARN of the Livy service account.	""
sparkServiceAccount.create	IF true, creates the Spark service account in <code>.Release.Namespace</code>	TRUE

Property	Description	Default
sparkServiceAccount.name	The name of the service account to use for Spark. If you don't set this property and create a Spark service account, Amazon EMR on EKS automatically generates a name with the <code>fullnameOverride</code> property with <code>-spark-livy</code> suffix.	"emr-containers-sa-spark-livy"
service.name	Name of the Livy service	"emr-containers-livy"
service.annotations	Livy service annotations	{}
loadbalancer.enabled	Whether to create a load balancer for the Livy service used to expose the Livy endpoint outside of the Amazon EKS cluster.	FALSE
loadbalancer.internal	Whether to configure the Livy endpoint as internal to the VPC or external. Setting this property to FALSE exposes the endpoint to sources outside of the VPC. We recommend securing your endpoint with TLS/SSL. For more information, see Setting up TLS and SSL encryption .	FALSE

Property	Description	Default
imagePullSecrets	The list of <code>imagePullSecret</code> names to use to pull Livy image from private repositories.	[]
resources	The resource requests and limits for Livy containers.	{}
nodeSelector	The nodes for which to schedule Livy pods.	{}
tolerations	A list containing the Livy pods tolerations to define.	[]
affinity	The Livy pods affinity rules.	{}
persistence.enabled	If true, enables persistence for sessions directories.	FALSE
persistence.subPath	The PVC subpath to mount to sessions directories.	""
persistence.existingClaim	The PVC to use instead of creating a new one.	{}

Property	Description	Default
<code>persistence.storageClass</code>	The storage class to use. To define this parameter, use the format <code>storageClassName: <storageClass></code> . Setting this parameter to "-" disables dynamic provisioning. If you set this parameter to null or don't specify anything, Amazon EMR on EKS doesn't set a <code>storageClassName</code> and uses the default provisioner.	""
<code>persistence.accessMode</code>	The PVC access mode.	<code>ReadWriteOnce</code>
<code>persistence.size</code>	The PVC size.	<code>20Gi</code>
<code>persistence.annotations</code>	Additional annotations for the PVC.	<code>{}</code>
<code>env.*</code>	Additional envs to set to Livy container. For more information, see Inputting your own Livy and Spark configurations while installing Livy .	<code>{}</code>
<code>envFrom.*</code>	Additional envs to set to Livy from a Kubernetes config map or secret.	<code>[]</code>
<code>livyConf.*</code>	Additional <code>livy.conf</code> entries to set from a mounted Kubernetes config map or secret.	<code>{}</code>

Property	Description	Default
<code>sparkDefaultsConf.*</code>	Additional <code>spark-defaults.conf</code> entries to set from a mounted Kubernetes config map or secret.	<code>{}</code>

Troubleshoot common environment-variable format errors

When you input Livy and Spark configurations, there are environment-variable formats that aren't supported and can cause errors. The procedure takes you through a series of steps to help ensure that you use correct formats.

Inputting your own Livy and Spark configurations while installing Livy

You can configure any Apache Livy or Apache Spark environment variable with the `env.*` Helm property. Follow the steps below to convert the example configuration `example.config.with-dash.withUppercase` to a supported environment variable format.

1. Replace uppercase letters with a 1 and a lowercase of the letter. For example, `example.config.with-dash.withUppercase` becomes `example.config.with-dash.with1uppercase`.
2. Replace dashes (-) with 0. For example, `example.config.with-dash.with1uppercase` becomes `example.config.with0dash.with1uppercase`.
3. Replace dots (.) with underscores (_). For example, `example.config.with0dash.with1uppercase` becomes `example_config_with0dash_with1uppercase`.
4. Replace all lowercase letters with uppercase letters.
5. Add the prefix `LIVY_` to the variable name.
6. Use the variable while installing Livy through the helm chart using the format `--set env.YOUR_VARIABLE_NAME.value=yourvalue`

For example, to set the Livy and Spark configurations `livy.server.recovery.state-store = filesystem` and `spark.kubernetes.executor.podNamePrefix = my-prefix`, use these Helm properties:

```
–set env.LIVY_LIVY_SERVER_RECOVERY_STATESTORE.value=filesystem  
–set env.LIVY_SPARK_KUBERNETES_EXECUTOR_PODNAMEPREFIX.value=myprefix
```

Managing Amazon EMR on EKS job runs

The following sections cover topics that help you manage your Amazon EMR on EKS job runs. These include configuring job run parameters when you use the AWS CLI, configuring how your log data is stored, running Spark SQL scripts to run queries, understanding job run states, and knowing how to monitor jobs. You can work through these topics, generally in order, if you want to set up and complete a job run to process data.

Topics

- [Managing job runs with the AWS CLI](#)
- [Running Spark SQL scripts through the StartJobRun API](#)
- [Job run states](#)
- [Viewing jobs in the Amazon EMR console](#)
- [Common errors when running jobs](#)

Managing job runs with the AWS CLI

This topic covers how to manage job runs with the AWS Command Line Interface (AWS CLI). It goes into detail regarding properties, like security parameters, the driver, and various override settings. It also includes subtopics that cover various ways to configure logging.

Topics

- [Options for configuring a job run](#)
- [Configure a job run to use Amazon S3 logs](#)
- [Configure a job run to use Amazon CloudWatch Logs](#)
- [List job runs](#)
- [Describe a job run](#)
- [Cancel a job run](#)

Options for configuring a job run

Use the following options to configure job run parameters:

- `--execution-role-arn`: You must provide an IAM role that is used for running jobs. For more information, see [Using job execution roles with Amazon EMR on EKS](#).
- `--release-label`: You can deploy Amazon EMR on EKS with Amazon EMR versions 5.32.0 and 6.2.0 and later. Amazon EMR on EKS is not supported in previous Amazon EMR release versions. For more information, see [Amazon EMR on EKS releases](#).
- `--job-driver`: Job driver is used to provide input on the main job. This is a union type field where you can only pass one of the values for the job type that you want to run. Supported job types include:
 - Spark submit jobs - Used to run a command through Spark submit. You can use this job type to run Scala, PySpark, SparkR, SparkSQL and any other supported jobs through Spark Submit. This job type has the following parameters:
 - Entrypoint - This is the HCFS (Hadoop compatible file system) reference to the main jar/py file you want to run.
 - EntryPointArguments - This is an array of arguments you want to pass to your main jar/py file. You should handle reading these parameters using your entrypoint code. Each argument in the array should be separated by a comma. EntryPointArguments cannot contain brackets or parentheses, such as `()`, `{}`, or `[]`.
 - SparkSubmitParameters - These are the additional spark parameters you want to send to the job. Use this parameter to override default Spark properties such as driver memory or number of executors like `—conf` or `—class`. For additional information, see [Launching Applications with spark-submit](#).
 - Spark SQL jobs - Used to run a SQL query file through Spark SQL. You can use this job type to run SparkSQL jobs. This job type has the following parameters:
 - Entrypoint - This is the HCFS (Hadoop compatible file system) reference to the SQL query file you want to run.

For a list of additional Spark parameters you can use for a Spark SQL job, see [Running Spark SQL scripts through the StartJobRun API](#).

- `--configuration-overrides`: You can override the default configurations for applications by supplying a configuration object. You can use a shorthand syntax to provide the configuration or you can reference the configuration object in a JSON file. Configuration objects consist of a classification, properties, and optional nested configurations. Properties consist of the settings

you want to override in that file. You can specify multiple classifications for multiple applications in a single JSON object. The configuration classifications that are available vary by Amazon EMR release version. For a list of configuration classifications that are available for each release version of Amazon EMR, see [Amazon EMR on EKS releases](#).

If you pass the same configuration in an application override and in Spark submit parameters, the Spark submit parameters take precedence. The complete configuration priority list follows, in order of highest priority to lowest priority.

- Configuration supplied when creating `SparkSession`.
- Configuration supplied as part of `sparkSubmitParameters` using `–conf`.
- Configuration provided as part of application overrides.
- Optimized configurations chosen by Amazon EMR for the release.
- Default open source configurations for the application.

To monitor job runs using Amazon CloudWatch or Amazon S3, you must provide the configuration details for CloudWatch. For more information, see [Configure a job run to use Amazon S3 logs](#) and [Configure a job run to use Amazon CloudWatch Logs](#). If the S3 bucket or CloudWatch log group does not exist, then Amazon EMR creates it before uploading logs to the bucket.

- For an additional list of Kubernetes configuration options, see [Spark Properties on Kubernetes](#).

The following Spark configurations are not supported.

- `spark.kubernetes.authenticate.driver.serviceAccountName`
- `spark.kubernetes.authenticate.executor.serviceAccountName`
- `spark.kubernetes.namespace`
- `spark.kubernetes.driver.pod.name`
- `spark.kubernetes.container.image.pullPolicy`
- `spark.kubernetes.container.image`

 Note

You can use `spark.kubernetes.container.image` for customized Docker images. For more information, see [Customizing Docker images for Amazon EMR on EKS](#).

Configure a job run to use Amazon S3 logs

To be able to monitor the job progress and to troubleshoot failures, you must configure your jobs to send log information to Amazon S3, Amazon CloudWatch Logs, or both. This topic helps you get started publishing application logs to Amazon S3 on your jobs that are launched with Amazon EMR on EKS.

S3 logs IAM policy

Before your jobs can send log data to Amazon S3, the following permissions must be included in the permissions policy for the job execution role. Replace *amzn-s3-demo-logging-bucket* with the name of your logging bucket.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::amzn-s3-demo-bucket",
        "arn:aws:s3:::amzn-s3-demo-bucket/*"
      ],
      "Sid": "AllowS3Putobject"
    }
  ]
}
```

Note

Amazon EMR on EKS can also create an Amazon S3 bucket. If an Amazon S3 bucket is not available, include the `s3:CreateBucket` permission in the IAM policy.

After you've given your execution role the proper permissions to send logs to Amazon S3, your log data are sent to the following Amazon S3 locations when `s3MonitoringConfiguration` is passed in the `monitoringConfiguration` section of a `start-job-run` request, as shown in [Managing job runs with the AWS CLI](#).

- Submitter Logs - `/logUri/virtual-cluster-id/jobs/job-id/containers/pod-name/(stderr.gz/stdout.gz)`
- Driver Logs - `/logUri/virtual-cluster-id/jobs/job-id/containers/spark-application-id/spark-job-id-driver/(stderr.gz/stdout.gz)`
- Executor Logs - `/logUri/virtual-cluster-id/jobs/job-id/containers/spark-application-id/executor-pod-name/(stderr.gz/stdout.gz)`

Configure a job run to use Amazon CloudWatch Logs

To monitor job progress and to troubleshoot failures, you must configure your jobs to send log information to Amazon S3, Amazon CloudWatch Logs, or both. This topic helps you get started using CloudWatch Logs on your jobs that are launched with Amazon EMR on EKS. For more information about CloudWatch Logs, see [Monitoring Log Files](#) in the Amazon CloudWatch User Guide.

CloudWatch Logs IAM policy

For your jobs to send log data to CloudWatch Logs, the following permissions must be included in the permissions policy for the job execution role. Replace `my_log_group_name` and `my_log_stream_prefix` with names of your CloudWatch log group and log stream names, respectively. Amazon EMR on EKS creates the log group and log stream if they do not exist as long as the execution role ARN has appropriate permissions.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "logs:CreateLogStream",
        "logs:DescribeLogGroups",
```

```

        "logs:DescribeLogStreams"
    ],
    "Resource": [
        "arn:aws:logs:*:*:*"
    ],
    "Sid": "AllowLOGSCreateLogstream"
},
{
    "Effect": "Allow",
    "Action": [
        "logs:PutLogEvents"
    ],
    "Resource": [
        "arn:aws:logs:*:*:log-group:my_log_group_name:log-
stream:my_log_stream_prefix/*"
    ],
    "Sid": "AllowLOGSPutLogEvents"
}
]
}

```

Note

Amazon EMR on EKS can also create a log stream. If a log stream does not exist, the IAM policy should include the "logs:CreateLogGroup" permission.

After you've given your execution role the proper permissions, your application sends its log data to CloudWatch Logs when `cloudWatchMonitoringConfiguration` is passed in the `monitoringConfiguration` section of a `start-job-run` request, as shown in [Managing job runs with the AWS CLI](#).

In the `StartJobRun` API, *log_group_name* is the log group name for CloudWatch, and *log_stream_prefix* is the log stream name prefix for CloudWatch. You can view and search these logs in the AWS Management Console.

- Submitter logs - *logGroup/logStreamPrefix/virtual-cluster-id/jobs/job-id/containers/pod-name/(stderr/stdout)*
- Driver logs - *logGroup/logStreamPrefix/virtual-cluster-id/jobs/job-id/containers/spark-application-id/spark-job-id-driver/(stderr/stdout)*

- Executor logs - *logGroup/logStreamPrefix/virtual-cluster-id/jobs/job-id/containers/spark-application-id/executor-pod-name/(stderr/stdout)*

List job runs

You can run `list-job-run` to show the states of job runs, as the following example demonstrates.

```
aws emr-containers list-job-runs --virtual-cluster-id <cluster-id>
```

Describe a job run

You can run `describe-job-run` to get more details about the job, such as job state, state details, and job name, as the following example demonstrates.

```
aws emr-containers describe-job-run --virtual-cluster-id cluster-id --id job-run-id
```

Cancel a job run

You can run `cancel-job-run` to cancel running jobs, as the following example demonstrates.

```
aws emr-containers cancel-job-run --virtual-cluster-id cluster-id --id job-run-id
```

Running Spark SQL scripts through the StartJobRun API

Amazon EMR on EKS releases 6.7.0 and higher include a Spark SQL job driver so that you can run Spark SQL scripts through the `StartJobRun` API. You can supply SQL entry-point files to directly run Spark SQL queries on Amazon EMR on EKS with the `StartJobRun` API, without any modifications to existing Spark SQL scripts. The following table lists Spark parameters that are supported for the Spark SQL jobs through the `StartJobRun` API.

You can choose from the following Spark parameters to send to a Spark SQL job. Use these parameters to override default Spark properties.

Option	Description
<code>--name NAME</code>	Application Name

Option	Description
<code>--jars JARS</code>	Comma separated list of jars to be included with driver and execute classpath.
<code>--packages</code>	Comma-separated list of maven coordinates of jars to include on the driver and executor classpaths.
<code>--exclude-packages</code>	Comma-separated list of groupId:artifactId, to exclude while resolving the dependencies provided in <code>--packages</code> to avoid dependency conflicts.
<code>--repositories</code>	Comma-separated list of additional remote repositories to search for the maven coordinates given with <code>--packages</code> .
<code>--files FILES</code>	Comma-separated list of files to be placed in the working directory of each executor.
<code>--conf PROP=VALUE</code>	Spark configuration property.
<code>--properties-file FILE</code>	Path to a file from which to load extra properties.
<code>--driver-memory MEM</code>	Memory for driver. Default 1024MB.
<code>--driver-java-options</code>	Extra Java options to pass to the driver.
<code>--driver-library-path</code>	Extra library path entries to pass to the driver.
<code>--driver-class-path</code>	Extra classpath entries to pass to the driver.
<code>--executor-memory MEM</code>	Memory per executor. Default 1GB.
<code>--driver-cores NUM</code>	Number of cores used by the driver.
<code>--total-executor-cores NUM</code>	Total cores for all executors.

Option	Description
<code>--executor-cores NUM</code>	Number of cores used by each executor.
<code>--num-executors NUM</code>	Number of executors to launch.
<code>-hivevar <key=value></code>	Variable substitution to apply to Hive commands, for example, <code>-hivevar A=B</code>
<code>-hiveconf <property=value></code>	Value to use for the given property.

For a Spark SQL job, create a `start-job-run-request.json` file and specify the required parameters for your job run, as in the following example:

```
{
  "name": "myjob",
  "virtualClusterId": "123456",
  "executionRoleArn": "iam_role_name_for_job_execution",
  "releaseLabel": "emr-6.7.0-latest",
  "jobDriver": {
    "sparkSqlJobDriver": {
      "entryPoint": "entryPoint_location",
      "sparkSqlParameters": "--conf spark.executor.instances=2 --conf
spark.executor.memory=2G --conf spark.executor.cores=2 --conf spark.driver.cores=1"
    }
  },
  "configurationOverrides": {
    "applicationConfiguration": [
      {
        "classification": "spark-defaults",
        "properties": {
          "spark.driver.memory": "2G"
        }
      }
    ]
  },
  "monitoringConfiguration": {
    "persistentAppUI": "ENABLED",
    "cloudWatchMonitoringConfiguration": {
      "logGroupName": "my_log_group",
      "logStreamNamePrefix": "log_stream_prefix"
    }
  },
}
```

```
    "s3MonitoringConfiguration": {  
      "logUri": "s3://my_s3_log_location"  
    }  
  }  
}  
}
```

Job run states

When you submit a job run to an Amazon EMR on EKS job queue, the job run enters the PENDING state. It then passes through the following states until it succeeds (exits with code 0) or fails (exits with a non-zero code).

Job runs can have the following states:

- **PENDING** - The initial job state when the job run is submitted to Amazon EMR on EKS. The job is waiting to be submitted to the virtual cluster, and Amazon EMR on EKS is working on submitting this job.
- **SUBMITTED** - A job run that has been successfully submitted to the virtual cluster. The cluster scheduler then tries to run this job on the cluster.
- **RUNNING** - A job run that is running in the virtual cluster. In Spark applications, this means that the Spark driver process is in the running state.
- **FAILED** - A job run that failed to be submitted to the virtual cluster or that completed unsuccessfully. Look at `StateDetails` and `FailureReason` to find additional information about this job failure.
- **COMPLETED** - A job run that has completed successfully.
- **CANCEL_PENDING** - A job run has been requested for cancellation. Amazon EMR on EKS is trying to cancel the job on the virtual cluster.
- **CANCELLED** - A job run that was cancelled successfully.

Viewing jobs in the Amazon EMR console

Job run data is available to view, so you can monitor each job as it passes through the states. To view jobs in the Amazon EMR console, perform the following steps.

1. In the Amazon EMR console lefthand menu, under Amazon EMR on EKS, choose **Virtual clusters**.
2. From the list of virtual clusters, select the virtual cluster for which you want to view jobs.

3. On the **Job runs** table, select **View logs** to view the details of a job run.

 Note

Support for the one-click experience is enabled by default. It can be turned off by setting `persistentAppUI` to `DISABLED` in `monitoringConfiguration` during job submission. For more information, see [View Persistent Application User Interfaces](#).

Common errors when running jobs

The following errors may occur when you run `StartJobRun` API. The table lists each error and provides mitigation steps so you can address issues quickly.

Error Message	Error Condition	Recommended Next Step
error: argument <i>--argument</i> is required	Required parameters are missing.	Add the missing arguments to the API request.
An error occurred (AccessDeniedException) when calling the StartJobRun operation: User: <i>ARN</i> is not authorized to perform: emr-containers:StartJobRun	Execution role is missing.	See Using Using job execution roles with Amazon EMR on EKS .
An error occurred (AccessDeniedException) when calling the StartJobRun operation: User: <i>ARN</i> is not authorized to perform: emr-containers:StartJobRun	Caller doesn't have permission to the execution role [valid / not valid format] via condition keys.	See Using job execution roles with Amazon EMR on EKS .
An error occurred (AccessDeniedException) when calling the StartJobRun operation: User: <i>ARN</i> is not authorized to	Job submitter and Execution role ARN are from different accounts.	Ensure that job submitter and execution role ARN are from the same AWS account.

Error Message	Error Condition	Recommended Next Step
perform: emr-containers:StartJobRun		
1 validation error detected: Value <i>Role</i> at 'executionRoleArn' failed to satisfy the ARN regular expression pattern: <code>^arn:(aws[a-zA-Z0-9-]*):iam::(\d{12})?:(role(\u002F) (\u002F\u0021-\u007F)+\u002F))(\w+=,.\u002D)+</code>	Caller has permissions for the execution role via condition keys, but the role does not satisfy the constraints of ARN format.	Provide the execution role following the ARN format. See Using job execution roles with Amazon EMR on EKS .
An error occurred (ResourceNotFoundException) when calling the StartJobRun operation: Virtual cluster <i>Virtual Cluster ID</i> doesn't exist.	Virtual cluster ID is not found.	Provide a virtual cluster ID registered with Amazon EMR on EKS.
An error occurred (ValidationException) when calling the StartJobRun operation: Virtual cluster state <i>state</i> is not valid to create resource JobRun.	Virtual cluster is not ready to execute job.	See Virtual cluster states .
An error occurred (ResourceNotFoundException) when calling the StartJobRun operation: Release <i>RELEASE</i> doesn't exist.	The release specified in job submission is incorrect.	See Amazon EMR on EKS releases .

Error Message	Error Condition	Recommended Next Step
An error occurred (AccessDeniedException) when calling the StartJobRun operation: User: <i>ARN</i> is not authorized to perform: emr-containers:StartJobRun on resource: <i>ARN</i> with an explicit deny. An error occurred (AccessDeniedException) when calling the StartJobRun operation: User: <i>ARN</i> is not authorized to perform: emr-containers:StartJobRun on resource: <i>ARN</i>	User is not authorized to call StartJobRun.	See Using job execution roles with Amazon EMR on EKS .
An error occurred (ValidationException) when calling the StartJobRun operation: configurationOverrides.monitoringConfiguration.s3MonitoringConfiguration.logUri failed to satisfy constraint : %s	S3 path URI syntax is not valid.	logUri should be in the format of s3://...

The following errors may occur when you run DescribeJobRun API before the job runs.

Error Message	Error Condition	Recommended Next Step
stateDetails: JobRun submission failed.	Parameters in StartJobRun are not valid.	See Amazon EMR on EKS releases .
Classification <i>classification</i> not supported.		

Error Message	Error Condition	Recommended Next Step
failureReason: VALIDATION_ERROR state: FAILED.		
stateDetails: Cluster <i>EKS Cluster ID</i> does not exist. failureReason: CLUSTER_UNAVAILABLE state: FAILED	The EKS cluster is not available.	Check if the EKS cluster exists and has the right permissions. For more information, see Setting up Amazon EMR on EKS .
stateDetails: Cluster <i>EKS Cluster ID</i> does not have sufficient permissions. failureReason: CLUSTER_UNAVAILABLE state: FAILED	Amazon EMR does not have permissions to access the EKS cluster.	Verify that permissions are set up for Amazon EMR on the registered namespace. For more information, see Setting up Amazon EMR on EKS .
stateDetails: Cluster <i>EKS Cluster ID</i> is currently not reachable. failureReason: CLUSTER_UNAVAILABLE state: FAILED	EKS cluster is not reachable.	Check if EKS Cluster exists and has the right permissions. For more information, see Setting up Amazon EMR on EKS .
stateDetails: JobRun submission failed due to an internal error. failureReason: INTERNAL_ERROR state: FAILED	An internal error has occurred with the EKS cluster.	N/A

Error Message	Error Condition	Recommended Next Step
stateDetails: Cluster <i>EKS Cluster ID</i> does not have sufficient resources. failureReason: USER_ERROR state: FAILED	There are insufficient resources in the EKS cluster to run the job.	Add more capacity to the EKS node group or set up EKS Autoscaler. For more information, see Cluster Autoscaler .

The following errors may occur when you run DescribeJobRun API after the job runs.

Error Message	Error Condition	Recommended Next Step
stateDetails: Trouble monitoring your JobRun. Cluster <i>EKS Cluster ID</i> does not exist. failureReason: CLUSTER_UNAVAILABLE state: FAILED	The EKS cluster does not exist.	Check if EKS Cluster exists and has the right permissions. For more information, see Setting up Amazon EMR on EKS .
stateDetails: Trouble monitoring your JobRun. Cluster <i>EKS Cluster ID</i> does not have sufficient permissions. failureReason: CLUSTER_UNAVAILABLE state: FAILED	Amazon EMR does not have permissions to access the EKS cluster.	Verify that permissions are set up for Amazon EMR on the registered namespace. For more information, see Setting up Amazon EMR on EKS .
stateDetails: Trouble monitoring your JobRun.	The EKS cluster is not reachable.	Check if EKS Cluster exists and has the right permissions.

Error Message	Error Condition	Recommended Next Step
Cluster <i>EKS Cluster ID</i> is currently not reachable. failureReason: CLUSTER_UNAVAILABLE state: FAILED		ns. For more information, see Setting up Amazon EMR on EKS .
stateDetails: Trouble monitoring your JobRun due to an internal error failureReason: INTERNAL_ERROR state: FAILED	An internal error has occurred and is preventing JobRun monitoring.	N/A

The following error may occur when a job cannot start and the job waits in the SUBMITTED state for 15 minutes. This can be caused by a lack of cluster resources.

Error Message	Error Condition	Recommended Next Step
cluster timeout	The job has been in the SUBMITTED state for 15 minutes or more.	You can override the default setting of 15 minutes for this parameter with the configuration override shown below.

Use the following configuration to change the cluster timeout setting to 30 minutes. Notice that you provide the new `job-start-timeout` value in seconds:

```
{
  "configurationOverrides": {
    "applicationConfiguration": [{
      "classification": "emr-containers-defaults",
      "properties": {
        "job-start-timeout": "1800"
      }
    }]
  }
}
```

```
    }  
  }]  
}
```

Using job templates

A job template stores values that can be shared across `StartJobRun` API invocations when starting a job run. It supports two use cases:

- To prevent repetitive recurring `StartJobRun` API request values.
- To enforce a rule that certain values must be provided via `StartJobRun` API requests.

Job templates enable you to define a reusable template for job runs to apply additional customization, for example:

- Configuring executor and driver compute capacity
- Setting security and governance properties such as IAM roles
- Customizing a docker image to use across multiple applications and data pipelines

The following topics provide detailed information on using templates, including how to use them to start a job run and how to change template parameters.

Topics

- [Create and using a job template to start a job run](#)
- [Defining job template parameters](#)
- [Controlling access to job templates](#)

Create and using a job template to start a job run

This section describes creating a job template and using the template to start a job run with the AWS Command Line Interface (AWS CLI).

To create a job template

1. Create a `create-job-template-request.json` file and specify the required parameters for your job template, as shown in the following example JSON file. For information about all available parameters, see the [CreateJobTemplate](#) API.

Most values that are required for the `StartJobRun` API are also required for `jobTemplateData`. If you want to use placeholders for any parameters and provide values when invoking `StartJobRun` using a job template, please see the next section on job template parameters.

```
{
  "name": "mytemplate",
  "jobTemplateData": {
    "executionRoleArn": "iam_role_arn_for_job_execution",
    "releaseLabel": "emr-6.7.0-latest",
    "jobDriver": {
      "sparkSubmitJobDriver": {
        "entryPoint": "entryPoint_location",
        "entryPointArguments": [ "argument1", "argument2", ... ],
        "sparkSubmitParameters": "--class <main_class> --conf
spark.executor.instances=2 --conf spark.executor.memory=2G --conf
spark.executor.cores=2 --conf spark.driver.cores=1"
      }
    },
    "configurationOverrides": {
      "applicationConfiguration": [
        {
          "classification": "spark-defaults",
          "properties": {
            "spark.driver.memory": "2G"
          }
        }
      ],
      "monitoringConfiguration": {
        "persistentAppUI": "ENABLED",
        "cloudWatchMonitoringConfiguration": {
          "logGroupName": "my_log_group",
          "logStreamNamePrefix": "log_stream_prefix"
        },
        "s3MonitoringConfiguration": {
          "logUri": "s3://my_s3_log_location/"
        }
      }
    }
  }
}
```

```
    }  
  }  
}
```

2. Use the `create-job-template` command with a path to the `create-job-template-request.json` file stored locally.

```
aws emr-containers create-job-template \  
--cli-input-json file://./create-job-template-request.json
```

To start a job run using a job template

Supply the virtual cluster id, job template id, and job name in the `StartJobRun` command, as shown in the following example.

```
aws emr-containers start-job-run \  
--virtual-cluster-id 123456 \  
--name myjob \  
--job-template-id 1234abcd
```

Defining job template parameters

Job template parameters allow you to specify variables in the job template. Values for these parameter variables will need to be specified when starting a job run using that job template. Job template parameters are specified in `${parameterName}` format. You can choose to specify any value in a `jobTemplateData` field as a job template parameter. For each of the job template parameter variables, specify its data type (`STRING` or `NUMBER`) and optionally a default value. The example below shows how you can specify job template parameters for entry point location, main class, and S3 log location values.

To specify entry point location, main class, and Amazon S3 log location as job template parameters

1. Create a `create-job-template-request.json` file and specify the required parameters for your job template, as shown in the following example JSON file. For more information about the parameters, see the [CreateJobTemplate](#) API.

```
{  
  "name": "mytemplate",
```



```

"jobTemplateData": {
  "executionRoleArn": "iam_role_arn_for_job_execution",
  "releaseLabel": "emr-6.7.0-latest",
  "jobDriver": {
    "sparkSubmitJobDriver": {
      "entryPoint": "${EntryPointLocation}",
      "entryPointArguments": [ "argument1", "argument2", ... ],
      "sparkSubmitParameters": "--class ${MainClass} --conf
spark.executor.instances=2 --conf spark.executor.memory=2G --conf
spark.executor.cores=2 --conf spark.driver.cores=1"
    }
  },
  "configurationOverrides": {
    "applicationConfiguration": [
      {
        "classification": "spark-defaults",
        "properties": {
          "spark.driver.memory": "2G"
        }
      }
    ],
    "monitoringConfiguration": {
      "persistentAppUI": "ENABLED",
      "cloudWatchMonitoringConfiguration": {
        "logGroupName": "my_log_group",
        "logStreamNamePrefix": "log_stream_prefix"
      },
      "s3MonitoringConfiguration": {
        "logUri": "${LogS3BucketUri}"
      }
    }
  },
  "parameterConfiguration": {
    "EntryPointLocation": {
      "type": "STRING"
    },
    "MainClass": {
      "type": "STRING",
      "defaultValue": "Main"
    },
    "LogS3BucketUri": {
      "type": "STRING",
      "defaultValue": "s3://my_s3_log_location/"
    }
  }
}

```

```
    }  
  }  
}
```

2. Use the `create-job-template` command with a path to the `create-job-template-request.json` file stored locally or in Amazon S3.

```
aws emr-containers create-job-template \  
--cli-input-json file://./create-job-template-request.json
```

To start a job run using job template with job template parameters

To start a job run with a job template containing job template parameters, specify the job template id as well as values for job template parameters in the `StartJobRun` API request as shown below.

```
aws emr-containers start-job-run \  
--virtual-cluster-id 123456 \  
--name myjob \  
--job-template-id 1234abcd \  
--job-template-parameters '{"EntryPointLocation": "entry_point_location", "MainClass":  
"ExampleMainClass", "LogS3BucketUri": "s3://example_s3_bucket/"}'
```

Controlling access to job templates

`StartJobRun` policy lets you enforce that a user or a role can only run jobs using job templates that you specify and cannot run `StartJobRun` operations without using the specified job templates. To achieve this, first ensure that you give the user or role a read permission to the specified job templates as shown below.

JSON

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "emr-containers:DescribeJobRun"  
      ],  
      "Resource": [  

```

```

        "arn:aws:emr-containers:*:*:jobtemplate/job_template_1_id",
        "arn:aws:emr-containers:*:*:jobtemplate/job_template_2_id"
    ],
    "Sid": "AllowEMRCONTAINERSDescribejobtemplate"
}
]
}

```

To enforce that a user or role is able to invoke StartJobRun operation only when using specified job templates, you can assign the following StartJobRun policy permission to a given user or role.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "emr-containers:StartJobRun"
      ],
      "Resource": [
        "arn:aws:emr-containers:*:*/virtualclusters/virtual_cluster_id"
      ],
      "Condition": {
        "ArnLike": {
          "emr-containers:JobTemplateArn": [
            "arn:aws:emr-containers:*:*:jobtemplate/job_template_1_id",
            "arn:aws:emr-containers:*:*:jobtemplate/job_template_2_id"
          ]
        }
      },
      "Sid": "AllowEMRCONTAINERSStartjobrun"
    }
  ]
}

```

If the job template specifies a job template parameter inside the execution role ARN field, then the user will be able to provide a value for this parameter and thus be able to invoke StartJobRun

using an arbitrary execution role. To restrict the execution roles the user can provide, see **Controlling access to the execution role** in [Using job execution roles with Amazon EMR on EKS](#).

If no condition is specified in the above `StartJobRun` action policy for a given user or a role, the user or the role will be allowed to invoke `StartJobRun` action on the specified virtual cluster using an arbitrary job template that they have read access to or using an arbitrary execution role.

Using pod templates

Beginning with Amazon EMR versions 5.33.0 or 6.3.0, Amazon EMR on EKS supports Spark's pod template feature. A pod is a group of one or more containers, with shared storage and network resources, and a specification for how to run the containers. Pod templates are specifications that determine how to run each pod. You can use pod template files to define the driver or executor pod's configurations that Spark configurations do not support. For more information about the Spark's pod template feature, see [Pod Template](#).

Note

The pod template feature only works with driver and executor pods. You cannot configure job submitter pods using the pod template.

Common scenarios

You can define how to run Spark jobs on shared EKS clusters by using pod templates with Amazon EMR on EKS and save costs and improve resource utilization and performance.

- To reduce costs, you can schedule Spark driver tasks to run on Amazon EC2 On-Demand Instances while scheduling Spark executor tasks to run on Amazon EC2 Spot Instances.
- To increase resource utilization, you can support multiple teams running their workloads on the same EKS cluster. Each team will get a designated Amazon EC2 node group to run their workloads on. You can use pod templates to apply a corresponding toleration to their workload.
- To improve monitoring, you can run a separate logging container to forward logs to your existing monitoring application.

For example, the following pod template file demonstrates a common usage scenario.

```
apiVersion: v1
```

```

kind: Pod
spec:
  volumes:
    - name: source-data-volume
      emptyDir: {}
    - name: metrics-files-volume
      emptyDir: {}
  nodeSelector:
    eks.amazonaws.com/nodegroup: emr-containers-nodegroup
  containers:
    - name: spark-kubernetes-driver # This will be interpreted as driver Spark main
      container
      env:
        - name: RANDOM
          value: "random"
      volumeMounts:
        - name: shared-volume
          mountPath: /var/data
        - name: metrics-files-volume
          mountPath: /var/metrics/data
    - name: custom-side-car-container # Sidecar container
      image: <side_car_container_image>
      env:
        - name: RANDOM_SIDE CAR
          value: random
      volumeMounts:
        - name: metrics-files-volume
          mountPath: /var/metrics/data
      command:
        - /bin/sh
        - '-c'
        - <command-to-upload-metrics-files>
  initContainers:
    - name: spark-init-container-driver # Init container
      image: <spark-pre-step-image>
      volumeMounts:
        - name: source-data-volume # Use EMR predefined volumes
          mountPath: /var/data
      command:
        - /bin/sh
        - '-c'
        - <command-to-download-dependency-jars>

```

The pod template completes the following tasks:

- Add a new [init container](#) that is executed before the Spark main container starts. The init container shares the [EmptyDir volume](#) called `source-data-volume` with the Spark main container. You can have your init container run initialization steps, such as downloading dependencies or generating input data. Then the Spark main container consumes the data.
- Add another [sidecar container](#) that is executed along with the Spark main container. The two containers are sharing another EmptyDir volume called `metrics-files-volume`. Your Spark job can generate metrics, such as Prometheus metrics. Then the Spark job can put the metrics into a file and have the sidecar container upload the files to your own BI system for future analysis.
- Add a new environment variable to the Spark main container. You can have your job consume the environment variable.
- Define a [node selector](#), so that the pod is only scheduled on the `emr-containers-nodegroup` node group. This helps to isolate compute resources across jobs and teams.

Enabling pod templates with Amazon EMR on EKS

To enable the pod template feature with Amazon EMR on EKS, configure the Spark properties `spark.kubernetes.driver.podTemplateFile` and `spark.kubernetes.executor.podTemplateFile` to point to the pod template files in Amazon S3. Spark then downloads the pod template file and uses it to construct driver and executor pods.

Note

Spark uses the job execution role to load the pod template, so the job execution role must have permissions to access Amazon S3 to load the pod templates. For more information, see [Create a job execution role](#).

You can use the `SparkSubmitParameters` to specify the Amazon S3 path to the pod template, as the following job run JSON file demonstrates.

```
{
  "name": "myjob",
  "virtualClusterId": "123456",
```

```

"executionRoleArn": "iam_role_name_for_job_execution",
"releaseLabel": "release_label",
"jobDriver": {
  "sparkSubmitJobDriver": {
    "entryPoint": "entryPoint_location",
    "entryPointArguments": [argument1, argument2, ...],
    "sparkSubmitParameters": "--class <main_class> \
      --conf
spark.kubernetes.driver.podTemplateFile=s3://path_to_driver_pod_template \
      --conf
spark.kubernetes.executor.podTemplateFile=s3://path_to_executor_pod_template \
      --conf spark.executor.instances=2 \
      --conf spark.executor.memory=2G \
      --conf spark.executor.cores=2 \
      --conf spark.driver.cores=1"
  }
}
}

```

Alternatively, you can use the `configurationOverrides` to specify the Amazon S3 path to the pod template, as the following job run JSON file demonstrates.

```

{
  "name": "myjob",
  "virtualClusterId": "123456",
  "executionRoleArn": "iam_role_name_for_job_execution",
  "releaseLabel": "release_label",
  "jobDriver": {
    "sparkSubmitJobDriver": {
      "entryPoint": "entryPoint_location",
      "entryPointArguments": [argument1, argument2, ...],
      "sparkSubmitParameters": "--class <main_class> \
        --conf spark.executor.instances=2 \
        --conf spark.executor.memory=2G \
        --conf spark.executor.cores=2 \
        --conf spark.driver.cores=1"
    }
  },
  "configurationOverrides": {
    "applicationConfiguration": [
      {
        "classification": "spark-defaults",
        "properties": {

```

```
        "spark.driver.memory": "2G",
        "spark.kubernetes.driver.podTemplateFile": "s3://path_to_driver_pod_template",

    "spark.kubernetes.executor.podTemplateFile": "s3://path_to_executor_pod_template"
    }
  }
]
}
}
```

Note

1. You need to follow the security guidelines when using the pod template feature with Amazon EMR on EKS, such as isolating untrusted application code. For more information, see [Amazon EMR on EKS security best practices](#).
2. You cannot change the Spark main container names by using `spark.kubernetes.driver.podTemplateContainerName` and `spark.kubernetes.executor.podTemplateContainerName`, because these names are hardcoded as `spark-kubernetes-driver` and `spark-kubernetes-executors`. If you want to customize the Spark main container, you must specify the container in a pod template with these hardcoded names.

Pod template fields

Consider the following field restrictions when configuring a pod template with Amazon EMR on EKS.

- Amazon EMR on EKS allows only the following fields in a pod template to enable proper job scheduling.

These are the allowed pod level fields:

- `apiVersion`
- `kind`
- `metadata`
- `spec.activeDeadlineSeconds`
- `spec.affinity`

- `spec.containers`
- `spec.enableServiceLinks`
- `spec.ephemeralContainers`
- `spec.hostAliases`
- `spec.hostname`
- `spec.imagePullSecrets`
- `spec.initContainers`
- `spec.nodeName`
- `spec.nodeSelector`
- `spec.overhead`
- `spec.preemptionPolicy`
- `spec.priority`
- `spec.priorityClassName`
- `spec.readinessGates`
- `spec.runtimeClassName`
- `spec.schedulerName`
- `spec.subdomain`
- `spec.terminationGracePeriodSeconds`
- `spec.tolerations`
- `spec.topologySpreadConstraints`
- `spec.volumes`

These are the allowed Spark main container level fields:

- `env`
- `envFrom`
- `name`
- `lifecycle`
- `livenessProbe`
- `readinessProbe`

- `resources`
- `startupProbe`

- `stdin`
- `stdinOnce`
- `terminationMessagePath`
- `terminationMessagePolicy`
- `tty`
- `volumeDevices`
- `volumeMounts`
- `workingDir`

When you use any disallowed fields in the pod template, Spark throws an exception and the job fails. The following example shows an error message in the Spark controller log due to disallowed fields.

```
Executor pod template validation failed.  
Field container.command in Spark main container not allowed but specified.
```

- Amazon EMR on EKS predefines the following parameters in a pod template. The fields that you specify in a pod template must not overlap with these fields.

These are the predefined volume names:

- `emr-container-communicate`
- `config-volume`
- `emr-container-application-log-dir`
- `emr-container-event-log-dir`
- `temp-data-dir`
- `mnt-dir`
- `home-dir`
- `emr-container-s3`

These are the predefined volume mounts that only apply to the Spark main container:

- Name: `emr-container-communicate`; MountPath: `/var/log/fluentd`
- Name: `emr-container-application-log-dir`; MountPath: `/var/log/spark/user`
- Name: `emr-container-event-log-dir`; MountPath: `/var/log/spark/apps`

- Name: `mnt-dir`; MountPath: `/mnt`

- Name: temp-data-dir; MountPath: /tmp
- Name: home-dir; MountPath: /home/hadoop

These are the predefined environment variables that only apply to the Spark main container:

- SPARK_CONTAINER_ID
- K8S_SPARK_LOG_URL_STDERR
- K8S_SPARK_LOG_URL_STDOUT
- SIDECAR_SIGNAL_FILE

 Note

You can still use these predefined volumes and mount them into your additional sidecar containers. For example, you can use `emr-container-application-log-dir` and mount it to your own sidecar container defined in the pod template.

If the fields you specify conflict with any of the predefined fields in the pod template, Spark throws an exception and the job fails. The following example shows an error message in the Spark application log due to conflicts with the predefined fields.

```
Defined volume mount path on main container must not overlap with reserved mount paths: [<reserved-paths>]
```

Sidecar container considerations

Amazon EMR controls the lifecycle of the pods provisioned by Amazon EMR on EKS. The sidecar containers should follow the same lifecycle of the Spark main container. If you inject additional sidecar containers into your pods, we recommend that you integrate with the pod lifecycle management that Amazon EMR defines so that the sidecar container can stop itself when the Spark main container exits.

To reduce costs, we recommend that you implement a process that prevents driver pods with sidecar containers from continuing to run after your job completes. The Spark driver deletes executor pods when the executor is done. However, when a driver program completes, the additional sidecar containers continue to run. The pod is billed until Amazon EMR on EKS cleans up the driver pod, usually less than one minute after the driver Spark main container completes. To

reduce costs, you can integrate your additional sidecar containers with the lifecycle management mechanism that Amazon EMR on EKS defines for both driver and executor pods, as described in the following section.

Spark main container in driver and executor pods sends heartbeat to a file `/var/log/fluentd/main-container-terminated` every two seconds. By adding the Amazon EMR predefined `emr-container-communicate` volume mount to your sidecar container, you can define a sub-process of your sidecar container to periodically track the last modified time for this file. The sub-process then stops itself if it discovers that the Spark main container stops the heartbeat for a longer duration.

The following example demonstrates a sub-process that tracks the heartbeat file and stops itself. Replace *`your_volume_mount`* with the path where you mount the predefined volume. The script is bundled inside the image used by sidecar container. In a pod template file, you can specify a sidecar container with the following commands `sub_process_script.sh` and `main_command`.

```
MOUNT_PATH="your_volume_mount"
FILE_TO_WATCH="$MOUNT_PATH/main-container-terminated"
INITIAL_HEARTBEAT_TIMEOUT_THRESHOLD=60
HEARTBEAT_TIMEOUT_THRESHOLD=15
SLEEP_DURATION=10

function terminate_main_process() {
    # Stop main process
}

# Waiting for the first heartbeat sent by Spark main container
echo "Waiting for file $FILE_TO_WATCH to appear..."
start_wait=$(date +%s)
while ! [[ -f "$FILE_TO_WATCH" ]]; do
    elapsed_wait=$(expr $(date +%s) - $start_wait)
    if [ "$elapsed_wait" -gt "$INITIAL_HEARTBEAT_TIMEOUT_THRESHOLD" ]; then
        echo "File $FILE_TO_WATCH not found after $INITIAL_HEARTBEAT_TIMEOUT_THRESHOLD
seconds; aborting"
        terminate_main_process
        exit 1
    fi
    sleep $SLEEP_DURATION;
done;
echo "Found file $FILE_TO_WATCH; watching for heartbeats..."

while [[ -f "$FILE_TO_WATCH" ]]; do
```

```

LAST_HEARTBEAT=$(stat -c %Y $FILE_TO_WATCH)
ELAPSED_TIME_SINCE_AFTER_HEARTBEAT=$(expr $(date +%s) - $LAST_HEARTBEAT)
if [ "$ELAPSED_TIME_SINCE_AFTER_HEARTBEAT" -gt "$HEARTBEAT_TIMEOUT_THRESHOLD" ];
then
    echo "Last heartbeat to file $FILE_TO_WATCH was more than
$HEARTBEAT_TIMEOUT_THRESHOLD seconds ago at $LAST_HEARTBEAT; terminating"
    terminate_main_process
    exit 0
fi
sleep $SLEEP_DURATION;
done;
echo "Outside of loop, main-container-terminated file no longer exists"

# The file will be deleted once the fluentd container is terminated

echo "The file $FILE_TO_WATCH doesn't exist any more;"
terminate_main_process
exit 0

```

Using job retry policies

In Amazon EMR on EKS versions 6.9.0 and later, you can set a retry policy for your job runs. Retry policies cause a job driver pod to be restarted automatically if it fails or is deleted. This makes long-running Spark streaming jobs more resilient to failures.

Setting a retry policy for a job

To configure a retry policy, you provide a `RetryPolicyConfiguration` field using the [StartJobRun](#) API. An example `retryPolicyConfiguration` is shown here:

```

aws emr-containers start-job-run \
--virtual-cluster-id cluster_id \
--name sample-job-name \
--execution-role-arn execution-role-arn \
--release-label emr-6.9.0-latest \
--job-driver '{
  "sparkSubmitJobDriver": {
    "entryPoint": "local:///usr/lib/spark/examples/src/main/python/pi.py",
    "entryPointArguments": [ "2" ],
    "sparkSubmitParameters": "--conf spark.executor.instances=2 --conf
spark.executor.memory=2G --conf spark.executor.cores=2 --conf spark.driver.cores=1"
  }
}'

```

```
}
}' \
--retry-policy-configuration '{
    "maxAttempts": 5
}' \
--configuration-overrides '{
    "monitoringConfiguration": {
        "cloudWatchMonitoringConfiguration": {
            "logGroupName": "my_log_group_name",
            "logStreamNamePrefix": "my_log_stream_prefix"
        },
        "s3MonitoringConfiguration": {
            "logUri": "s3://amzn-s3-demo-logging-bucket"
        }
    }
}'
```

Note

retryPolicyConfiguration is only available from AWS CLI 1.27.68 version onwards. To update the AWS CLI to the latest version, see [Installing or updating the latest version of the AWS CLI](#)

Configure the maxAttempts field with the maximum number of times you want the job driver pod to be restarted if it fails or is deleted. The execution interval between two job driver retry attempts is an exponential retry interval of (10 seconds, 20 seconds, 40 seconds ...) which is capped at 6 minutes, as described in the [Kubernetes documentation](#).

Note

Every additional job driver execution will be billed as another job run, and will be subject to [Amazon EMR on EKS pricing](#).

Retry policy configuration values

- **Default retry policy for a job:** StartJobRun includes a retry policy set to 1 maximum attempt by default. You can configure the retry policy as desired.

Note

If `maxAttempts` of the `retryPolicyConfiguration` is set to 1, it means that no retries will be done to bring up the driver pod on failure.

- **Disabling retry policy for a job:** To disable a retry policy, set the max attempts value in `retryPolicyConfiguration` to 1.

```
"retryPolicyConfiguration": {  
  "maxAttempts": 1  
}
```

- **Set `maxAttempts` for a job within the valid range:** `StartJobRun` call will fail if the `maxAttempts` value is outside the valid range. The valid `maxAttempts` range is from 1 to 2,147,483,647 (32-bit integer), the range supported for Kubernetes' `backOffLimit` configuration setting. For more information, see [Pod backoff failure policy](#) in the Kubernetes documentation. If the `maxAttempts` value is invalid, the following error message is returned:

```
{  
  "message": "Retry policy configuration's parameter value of maxAttempts is invalid"  
}
```

Retrieving a retry policy status for a job

You can view the status of the retry attempts for a job with the [ListJobRuns](#) and [DescribeJobRun](#) APIs. Once you request a job with an enabled retry policy configuration, the `ListJobRun` and `DescribeJobRun` responses will contain the status of the retry policy in the `RetryPolicyExecution` field. In addition, the `DescribeJobRun` response will contain the `RetryPolicyConfiguration` that was input in the `StartJobRun` request for the job.

Sample responses

ListJobRuns response

```
{  
  "jobRuns": [  
    ...  
    ...  
  ]  
}
```

```

    "retryPolicyExecution" : {
      "currentAttemptCount": 2
    }
    ...
    ...
  ]
}

```

DescribeJobRun response

```

{
  ...
  ...
  "retryPolicyConfiguration": {
    "maxAttempts": 5
  },
  "retryPolicyExecution" : {
    "currentAttemptCount": 2
  },
  ...
  ...
}

```

These fields will not be visible when retry policy is disabled in the job, as described below in [Retry policy configuration values](#).

Monitoring a job with a retry policy

When you enable a retry policy, a CloudWatch event is generated for every job driver that is created. To subscribe to these events, set up a CloudWatch event rule using the following command:

```

aws events put-rule \
--name cwe-test \
--event-pattern '{"detail-type": ["EMR Job Run New Driver Attempt"]}'

```

The event will return information on the `newDriverPodName`, `newDriverCreatedAt` timestamp, `previousDriverFailureMessage`, and `currentAttemptCount` of the job drivers. These events will not be created if the retry policy is disabled.

For more information on how to monitor your job with CloudWatch events, see [Monitor jobs with Amazon CloudWatch Events](#).

Finding logs for drivers and executors

Driver pod names follow the format `spark-<job id>-driver-<random-suffix>`. The same `random-suffix` is added to the executor pod names that the driver spawns. When you use this `random-suffix`, you can find logs for a driver and its associated executors. The `random-suffix` is only present if the [retry policy is enabled](#) for the job; otherwise, the `random-suffix` is absent.

For more information on how to configure jobs with monitoring configuration for logging, see [Run a Spark application](#).

Using Spark event log rotation

With Amazon EMR 6.3.0 and later, you can turn on the Spark event log rotation feature for Amazon EMR on EKS. Instead of generating a single event log file, this feature rotates the file based on your configured time interval and removes the oldest event log files.

Rotating Spark event logs can help you avoid potential issues with a large Spark event log file generated for long running or streaming jobs. For example, you start a long running Spark job with an event log enabled with the `persistentAppUI` parameter. The Spark driver generates an event log file. If the job runs for hours or days and there is a limited disk space on the Kubernetes node, the event log file can consume all available disk space. Turning on the Spark event log rotation feature solves the problem by splitting the log file into multiple files and removing the oldest files.

Note

This feature only works with Amazon EMR on EKS. Amazon EMR running on Amazon EC2 doesn't support Spark event log rotation.

To turn on the Spark event log rotation feature, configure the following Spark parameters:

- `spark.eventLog.rotation.enabled` - turns on log rotation. It is disabled by default in the Spark configuration file. Set it to true to turn on this feature.
- `spark.eventLog.rotation.interval` - specifies time interval for the log rotation. The minimum value is 60 seconds. The default value is 300 seconds.

- `spark.eventLog.rotation.minFileSize` - specifies a minimum file size to rotate the log file. The minimum and default value is 1 MB.
- `spark.eventLog.rotation.maxFilesToRetain` - specifies how many rotated log files to keep during cleanup. The valid range is 1 to 10. The default value is 2.

You can specify these parameters in the `sparkSubmitParameters` section of the [StartJobRun](#) API, as the following example shows.

```
"sparkSubmitParameters": "--class org.apache.spark.examples.SparkPi --conf  
  spark.eventLog.rotation.enabled=true --conf spark.eventLog.rotation.interval=300 --  
conf spark.eventLog.rotation.minFileSize=1m --conf  
  spark.eventLog.rotation.maxFilesToRetain=2"
```

Using Spark container log rotation

With Amazon EMR 6.11.0 and later, you can turn on the Spark container log rotation feature for Amazon EMR on EKS. Instead of generating a single `stdout` or `stderr` log file, this feature rotates the file based on your configured rotation size and removes the oldest log files from the container.

Rotating Spark container logs can help you avoid potential issues with a large Spark log files generated for long-running or streaming jobs. For example, you might start a long-running Spark job, and the Spark driver generates a container log file. If the job runs for hours or days and there is limited disk space on the Kubernetes node, the container log file can consume all available disk space. When you turn on Spark container log rotation, you split the log file into multiple files, and remove the oldest files.

To turn on the Spark container log rotation feature, configure the following Spark parameters:

containerLogRotationConfiguration

Include this parameter in `monitoringConfiguration` to turn on log rotation. It is disabled by default. You must use `containerLogRotationConfiguration` in addition to `s3MonitoringConfiguration`.

rotationSize

The `rotationSize` parameter specifies file size for the log rotation. The range of possible values is from 2KB to 2GB. The numeric unit portion of the `rotationSize` parameter is passed

as an integer. Since decimal values aren't supported, you can specify a rotation size of 1.5GB, for example, with the value 1500MB.

maxFilesToKeep

The `maxFilesToKeep` parameter specifies the maximum number of files to retain in container after rotation has taken place. The minimum value is 1, and the maximum value is 50.

You can specify these parameters in the `monitoringConfiguration` section of the `StartJobRun` API, as the following example shows. In this example, with `rotationSize = "10 MB"` and `maxFilesToKeep = 3`, Amazon EMR on EKS rotates your logs at 10 MB, generates a new log file, and then purges the oldest log file once the number of log files reaches 3.

```
{
  "name": "my-long-running-job",
  "virtualClusterId": "123456",
  "executionRoleArn": "iam_role_name_for_job_execution",
  "releaseLabel": "emr-6.11.0-latest",
  "jobDriver": {
    "sparkSubmitJobDriver": {
      "entryPoint": "entryPoint_location",
      "entryPointArguments": ["argument1", "argument2", ...],
      "sparkSubmitParameters": "--class main_class --conf spark.executor.instances=2
--conf spark.executor.memory=2G --conf spark.executor.cores=2 --conf
spark.driver.cores=1"
    }
  },
  "configurationOverrides": {
    "applicationConfiguration": [
      {
        "classification": "spark-defaults",
        "properties": {
          "spark.driver.memory": "2G"
        }
      }
    ]
  },
  "monitoringConfiguration": {
    "persistentAppUI": "ENABLED",
    "cloudWatchMonitoringConfiguration": {
      "logGroupName": "my_log_group",
      "logStreamNamePrefix": "log_stream_prefix"
    }
  },
}
```

```
{
  "s3MonitoringConfiguration": {
    "logUri": "s3://my_s3_log_location"
  },
  "containerLogRotationConfiguration": {
    "rotationSize": "10MB",
    "maxFilesToKeep": "3"
  }
}
```

To start a job run with Spark container log rotation, include a path to the json file that you configured with these parameters in the [StartJobRun](#) command.

```
aws emr-containers start-job-run \
--cli-input-json file://path-to-json-request-file
```

Using vertical autoscaling with Amazon EMR Spark jobs

Amazon EMR on EKS vertical autoscaling automatically tunes memory and CPU resources to adapt to the needs of the workload that you provide for Amazon EMR Spark applications. This simplifies resource management.

To track the real-time and historic resource utilization of your Amazon EMR Spark applications, vertical autoscaling leverages the Kubernetes [Vertical Pod Autoscaler \(VPA\)](#). The vertical autoscaling capability uses the data that VPA collects to automatically tune the memory and CPU resources assigned to your Spark applications. This simplified process enhances reliability and optimizes cost.

Topics

- [Setting up vertical autoscaling for Amazon EMR on EKS](#)
- [Getting started with vertical autoscaling for Amazon EMR on EKS](#)
- [Configuring vertical autoscaling for Amazon EMR on EKS](#)
- [Monitoring vertical autoscaling for Amazon EMR on EKS](#)
- [Uninstall the Amazon EMR on EKS vertical autoscaling operator](#)

Setting up vertical autoscaling for Amazon EMR on EKS

This topic helps you get your Amazon EKS cluster ready to submit Amazon EMR Spark jobs with vertical autoscaling. The setup process requires you to confirm or complete the tasks in the following sections:

Topics

- [Prerequisites](#)
- [Install the Operator Lifecycle Manager \(OLM\) on your Amazon EKS cluster](#)
- [Install the Amazon EMR on EKS vertical autoscaling operator](#)

Prerequisites

Complete the following tasks before you install the vertical autoscaling Kubernetes operator on your cluster. If you've already completed any of the prerequisites, you can skip those and move on to the next one.

- [Install or update to the latest version of the AWS CLI](#) – If you've already installed the AWS CLI, confirm that you have the latest version.
- [Install kubectl](#) – kubectl is a command line tool that you use to communicate with the Kubernetes API server. You need kubectl to install and monitor vertical autoscaling-related artifacts on your Amazon EKS cluster.
- [Install the Operator SDK](#) – Amazon EMR on EKS uses the Operator SDK as a package manager for the life of the vertical autoscaling operator that you install on your cluster.
- [Install Docker](#) – You need access to the Docker CLI to authenticate and fetch the vertical autoscaling-related Docker images to install on your Amazon EKS cluster.
- [Install the Kubernetes Metrics server](#) – You must first install metrics server so the vertical pod autoscaler can fetch metrics from the Kubernetes API server.
- [Get started with Amazon EKS – eksctl \(version 1.24 or higher\)](#) – Vertical autoscaling is supported with Amazon EKS versions 1.24 and higher. Once you create the cluster, [register it for use with Amazon EMR](#).
- [Select an Amazon EMR base image URI \(release 6.10.0 or higher\)](#) – Vertical autoscaling is supported with Amazon EMR releases 6.10.0 and higher.

Install the Operator Lifecycle Manager (OLM) on your Amazon EKS cluster

Use the Operator SDK CLI to install the Operator Lifecycle Manager (OLM) on the Amazon EMR on EKS cluster where you want to set up vertical autoscaling, as shown in the following example. Once you set it up, you can use OLM to install and manage the lifecycle of the [Amazon EMR vertical autoscaling operator](#).

```
operator-sdk olm install
```

To validate installation, run the `olm status` command:

```
operator-sdk olm status
```

Verify that the command returns a successful result, similar to the following example output:

```
INFO[0007] Successfully got OLM status for version X.XX
```

If your installation doesn't succeed, see [Troubleshooting Amazon EMR on EKS vertical autoscaling](#).

Install the Amazon EMR on EKS vertical autoscaling operator

Use the following steps to install the vertical autoscaling operator on your Amazon EKS cluster:

1. Set up the following environment variables that you will use to complete the installation:
 - **\$REGION** points to the AWS Region for your cluster. For example, `us-west-2`.
 - **\$ACCOUNT_ID** points to the Amazon ECR account ID for your Region. For more information, see [Amazon ECR registry accounts by Region](#).
 - **\$RELEASE** points to the Amazon EMR release that you want to use for your cluster. With vertical autoscaling, you must use Amazon EMR release 6.10.0 or higher.
2. Next, get authentication tokens to the [Amazon ECR registry](#) for the operator.

```
aws ecr get-login-password \
  --region region-id | docker login \
  --username AWS \
  --password-stdin $ACCOUNT_ID.dkr.ecr.region-id.amazonaws.com
```

3. Install the Amazon EMR on EKS vertical autoscaling operator with the following command:

```
ECR_URL=$ACCOUNT_ID.dkr.ecr.$REGION.amazonaws.com && \
```

```
REPO_DEST=dynamic-sizing-k8s-operator-olm-bundle && \  
BUNDLE_IMG=emr-$RELEASE-dynamic-sizing-k8s-operator && \  
operator-sdk run bundle \  
$ECR_URL/$REPO_DEST/$BUNDLE_IMG\:latest
```

This will create a release of the vertical autoscaling operator in the default namespace of your Amazon EKS cluster. Use this command to install in a different namespace:

```
operator-sdk run bundle \  
$ACCOUNT_ID.dkr.ecr.$REGION.amazonaws.com/dynamic-sizing-k8s-operator-olm-bundle/  
emr-$RELEASE-dynamic-sizing-k8s-operator\:latest \  
-n operator-namespace
```

Note

If the namespace that you specify doesn't exist, OLM won't install the operator. For more information, see [Kubernetes namespace not found](#).

4. Verify that you successfully installed the operator with the `kubectl` Kubernetes command-line tool.

```
kubectl get csv -n operator-namespace
```

The `kubectl` command should return your newly-deployed vertical autoscaler operator with a **Phase** status of **Succeeded**. If you've trouble with installation or setup, see [Troubleshooting Amazon EMR on EKS vertical autoscaling](#).

Getting started with vertical autoscaling for Amazon EMR on EKS

Use vertical autoscaling for Amazon EMR on EKS when you want automatic tuning of memory and CPU resources to adapt to your Amazon EMR Spark application workload. For more information, see [Using vertical autoscaling with Amazon EMR Spark jobs](#).

Submitting a Spark job with vertical autoscaling

When you submit a job through the [StartJobRun](#) API, add the following two configurations to the driver for your Spark job to turn on vertical autoscaling:

```
"spark.kubernetes.driver.annotation.emr-containers.amazonaws.com/  
dynamic.sizing":"true",  
"spark.kubernetes.driver.annotation.emr-containers.amazonaws.com/  
dynamic.sizing.signature":"YOUR_JOB_SIGNATURE"
```

In the code above, the first line enables the vertical autoscaling capability. The next line is a required signature configuration that lets you choose a signature for your job.

For more information on these configurations and acceptable parameter values, see [Configuring vertical autoscaling for Amazon EMR on EKS](#). By default, your job submits in the monitoring-only **Off** mode of vertical autoscaling. This monitoring state lets you compute and view resource recommendations without performing autoscaling. For more information, see [Vertical autoscaling modes](#).

The following example shows how to complete a sample `start-job-run` command with vertical autoscaling:

```
aws emr-containers start-job-run \  
--virtual-cluster-id $VIRTUAL_CLUSTER_ID \  
--name $JOB_NAME \  
--execution-role-arn $EMR_ROLE_ARN \  
--release-label emr-6.10.0-latest \  
--job-driver '{  
  "sparkSubmitJobDriver": {  
    "entryPoint": "local:///usr/lib/spark/examples/src/main/python/pi.py"  
  }  
' \  
--configuration-overrides '{  
  "applicationConfiguration": [{  
    "classification": "spark-defaults",  
    "properties": {  
      "spark.kubernetes.driver.annotation.emr-containers.amazonaws.com/  
dynamic.sizing": "true",  
      "spark.kubernetes.driver.annotation.emr-containers.amazonaws.com/  
dynamic.sizing.signature": "test-signature"  
    }  
  }]  
'
```


Verifying the vertical autoscaling functionality

To verify that vertical autoscaling works correctly for the submitted job, use `kubectl` to get the `verticalpodautoscaler` custom resource and view your scaling recommendations. For example, the following command queries for recommendations on the example job from the [Submitting a Spark job with vertical autoscaling](#) section:

```
kubectl get verticalpodautoscalers --all-namespaces \
-l=emr-containers.amazonaws.com/dynamic.sizing.signature=test-signature
```

The output from this query should resemble the following:

NAME	MODE	CPU	MEM
PROVIDED AGE			
ds-jceyefkxnhrvdzw6djum3naf2abm6o63a6dvjkkedqtkhlrf25eq-vpa	Off	3304504865	True
87m			

If your output doesn't look similar or contains an error code, see [Troubleshooting Amazon EMR on EKS vertical autoscaling](#) for steps to help resolve the issue.

Configuring vertical autoscaling for Amazon EMR on EKS

You can configure vertical autoscaling when you submit Amazon EMR Spark jobs through the [StartJobRun](#) API. Set the autoscaling-related configuration parameters on the Spark driver pod as shown in the example in [Submitting a Spark job with vertical autoscaling](#).

The Amazon EMR on EKS vertical autoscaling operator listens to driver pods that have autoscaling, then sets up integration with the Kubernetes Vertical Pod Autoscaler (VPA) with the settings on the driver pod. This facilitates resource tracking and autoscaling of Spark executor pods.

The following sections describe the parameters that you can use when you configure vertical autoscaling for your Amazon EKS cluster.

Note

Configure the feature toggle parameter as a label, and configure the remaining parameters as annotations on the Spark driver pod. The autoscaling parameters belong to the `emr-containers.amazonaws.com/` domain and have the `dynamic.sizing` prefix.

Required parameters

You must include the following two parameters on the Spark job driver when you submit your job:

Key	Description	Accepted values	Default value	Type	Spark parameter ¹
<code>dynamic.sizing</code>	Feature toggle	true, false	not set	label	<code>spark.kubernetes.driver.label.emr-containers.amazonaws.com/dynamic.sizing</code>
<code>dynamic.sizing.signature</code>	Job signature	<i>string</i>	not set	annotation	<code>spark.kubernetes.driver.annotation.emr-containers.amazonaws.com/dynamic.sizing.signature</code>

¹ Use this parameter as a `SparkSubmitParameter` or `ConfigurationOverride` in the `StartJobRun` API.

- **dynamic.sizing** – You can turn vertical autoscaling on and off with the `dynamic.sizing` label. To turn on vertical autoscaling, set `dynamic.sizing` to `true` on the Spark driver pod. If you omit this label or set it to any value other than `true`, vertical autoscaling is off.
- **dynamic.sizing.signature** – Set the job signature with the `dynamic.sizing.signature` annotation on the driver pod. Vertical autoscaling aggregates your resource usage data across

different runs of Amazon EMR Spark jobs to derive resource recommendations. You provide the unique identifier to tie the jobs together.

 Note

If your job recurs at a fixed interval such as daily or weekly, then your job signature should remain the same for each new instance of the job. This ensures that vertical autoscaling can compute and aggregate recommendations across different runs of the job.

¹ Use this parameter as a `SparkSubmitParameter` or `ConfigurationOverride` in the `StartJobRun` API.

Optional parameters

Vertical autoscaling also supports the following optional parameters. Set them as annotations on the driver pod.

Key	Description	Accepted values	Default value	Type	Spark parameter ¹
dynamic.sizing.mode	Vertical autoscaling mode	Off, Initial, Auto	Off	annotation	<code>spark.kubernetes.driver.annotation.emr-containers.amazonaws.com/dynamic.sizing.mode</code>
dynamic.sizing.scale.memory	Enables memory scaling	<i>true, false</i>	true	annotation	<code>spark.kubernetes.driver.annotation.</code>

Key	Description	Accepted values	Default value	Type	Spark parameter ¹
					mr-containers.amazonaws.com/dynamic.sizing.scale.memory
dynamic.sizing.scale.cpu	Turn CPU scaling on or off	<i>true, false</i>	false	annotation	spark.kubernetes.driver.annotation.emr-containers.amazonaws.com/dynamic.sizing.scale.cpu
dynamic.sizing.scale.memory.min	Minimum limit for memory scaling	<i>string, K8s resource quantity ex: 1G</i>	not set	annotation	spark.kubernetes.driver.annotation.emr-containers.amazonaws.com/dynamic.sizing.scale.memory.min

Key	Description	Accepted values	Default value	Type	Spark parameter ¹
dynamic.sizing.scale.memory.max	Maximum limit for memory scaling	<i>string, K8s resource quantity ex: 4G</i>	not set	annotation	spark.kubernetes.driver.annotation.emr-containers.amazonaws.com/dynamic.sizing.scale.memory.max
dynamic.sizing.scale.cpu.min	Minimum limit for CPU scaling	<i>string, K8s resource quantity ex: 1</i>	not set	annotation	spark.kubernetes.driver.annotation.emr-containers.amazonaws.com/dynamic.sizing.scale.cpu.min

Key	Description	Accepted values	Default value	Type	Spark parameter ¹
dynamic.sizing.scale.cpu.max	Maximum limit for CPU scaling	<i>string, K8s resource quantity ex: 2</i>	not set	annotation	<code>spark.kubernetes.driver.annotation.emr-containers.amazonaws.com/dynamic.sizing.scale.cpu.max</code>

Vertical autoscaling modes

The mode parameter maps to the different autoscaling modes that the VPA supports. Use the `dynamic.sizing.mode` annotation on the driver pod to set the mode. The following values are supported for this parameter:

- **Off** – A dry-run mode where you can monitor recommendations, but autoscaling is not performed. This is the default mode for vertical autoscaling. In this mode, the associated vertical pod autoscaler resource computes recommendations, and you can monitor the recommendations through tools like `kubectl`, Prometheus, and Grafana.
- **Initial** – In this mode, VPA autoscales resources when the job starts if recommendations are available based on historic runs of the job, such as in the case of a recurring job.
- **Auto** – In this mode, VPA evicts Spark executor pods, and autoscales them with the recommended resource settings when the Spark driver pod restarts them. Sometimes, the VPA evicts running Spark executor pods, so it might result in additional latency when it retries the interrupted executor.

Resource scaling

When you set up vertical autoscaling, you can choose whether to scale CPU and memory resources. Set the `dynamic.sizing.scale.cpu` and `dynamic.sizing.scale.memory` annotations to `true` or `false`. By default, CPU scaling is set to `false`, and memory scaling is set to `true`.

Resource minimums and maximums (Bounds)

Optionally, you can also set boundaries on the CPU and memory resources. Choose a minimum and maximum value for these resources with the `dynamic.sizing.[memory/cpu].[min/max]` annotations when you enable autoscaling. By default, the resources have no limitations. Set the annotations as string values that represent a Kubernetes resource quantity. For example, set `dynamic.sizing.memory.max` to `4G` to represent 4 GB.

Monitoring vertical autoscaling for Amazon EMR on EKS

You can use the **kubectl** Kubernetes command line tool to list the active, vertical autoscaling-related recommendations on your cluster. You can also view your tracked job signatures, and purge any unneeded resources that are associated with the signatures.

List the vertical autoscaling recommendations for your cluster

Use `kubectl` to get the `verticalpodautoscalers` resource, and view the current status and recommendations. The following example query returns all active resources on your Amazon EKS cluster.

```
kubectl get verticalpodautoscalers \
-o custom-columns="NAME:.metadata.name, \"\
\"SIGNATURE:.metadata.labels.emr-containers\\.amazonaws\\.com/dynamic\\.sizing
\\.signature, \"\
\"MODE:.spec.updatePolicy.updateMode, \"\
\"MEM:.status.recommendation.containerRecommendations[0].target.memory\" \
--all-namespaces
```

The output from this query resembles the following:

NAME	SIGNATURE	MODE	MEM
ds- <i>example-id-1</i> -vpa	<i>job-signature-1</i>	Off	<i>none</i>
ds- <i>example-id-2</i> -vpa	<i>job-signature-2</i>	Initial	12936384283

Query and delete the vertical autoscaling recommendations for your cluster

When you delete an Amazon EMR vertical autoscaling job-run resource, it automatically deletes the associated VPA object that tracks and stores recommendations.

The following example uses `kubectl` to purge recommendations for a job that is identified by a signature:

```
kubectl delete jobrun -n emr -l=emr-containers\.amazonaws\.com/dynamic\.sizing  
\.signature=integ-test  
jobrun.dynamicsizing.emr.services.k8s.aws "ds-job-signature" deleted
```

If you don't know the specific job signature, or want to purge all of the resources on the cluster, you can use `--all` or `--all-namespaces` in your command instead of the unique job ID, as shown in the following example:

```
kubectl delete jobruns --all --all-namespaces  
jobrun.dynamicsizing.emr.services.k8s.aws "ds-example-id" deleted
```

Uninstall the Amazon EMR on EKS vertical autoscaling operator

If you want to remove the vertical autoscaling operator from your Amazon EKS cluster, use the `cleanup` command with the Operator SDK CLI as shown in the following example. This also deletes upstream dependencies that installed with the operator, such as the Vertical Pod Autoscaler.

```
operator-sdk cleanup emr-dynamic-sizing
```

If there are any running jobs on the cluster when you delete the operator, those jobs continue to run without vertical autoscaling. If you submit jobs on the cluster after you delete the operator, Amazon EMR on EKS will ignore any vertical autoscaling-related parameters that you may have defined during [configuration](#).

Running interactive workloads on Amazon EMR on EKS

An *interactive endpoint* is a gateway that connects Amazon EMR Studio to Amazon EMR on EKS so that you can run interactive workloads. You can use interactive endpoints with EMR Studio to run interactive analytics with datasets in data stores like [Amazon S3](#) and [Amazon DynamoDB](#).

Use cases

- Create an ETL script with the EMR Studio IDE experience. The IDE ingests on-premises data and stores it in Amazon S3 after transformations for subsequent analysis.
- Use notebooks to explore datasets and train a machine-learning model to detect anomalies in the datasets.
- Create scripts that generate daily reports for analytic applications like business dashboards.

Topics

- [Overview of interactive endpoints](#)
- [Prerequisites to create an interactive endpoint on Amazon EMR on EKS](#)
- [Creating an interactive endpoint for your virtual cluster](#)
- [Configuring settings for interactive endpoints](#)
- [Monitoring interactive endpoints](#)
- [Using self-hosted Jupyter notebooks](#)
- [Getting information about interactive endpoints with CLI commands](#)

Overview of interactive endpoints

An *interactive endpoint* provides the capability for interactive clients like Amazon EMR Studio to connect to Amazon EMR on EKS clusters to run interactive workloads. The interactive endpoint is backed by a Jupyter Enterprise Gateway that provides the remote kernel lifecycle management capability that interactive clients need. *Kernels* are language-specific processes that interact with the Jupyter-based Amazon EMR Studio client to run interactive workloads.

Interactive endpoints support the following kernels:

- Python 3
- PySpark on Kubernetes

- Apache Spark with Scala

Note

Amazon EMR on EKS pricing applies for the interactive endpoints and kernels. For more information, see the [Amazon EMR on EKS pricing page](#).

The following entities are required for EMR Studio to connect with Amazon EMR on EKS.


- **Amazon EMR on EKS virtual cluster** – A *virtual cluster* is a Kubernetes namespace that you register Amazon EMR with. Amazon EMR uses virtual clusters to run jobs and host endpoints. You can back multiple virtual clusters with the same physical cluster. However, each virtual cluster maps to one namespace on an Amazon EKS cluster. Virtual clusters don't create any active resources that contribute to your bill or that require lifecycle management outside the service.
- **Amazon EMR on EKS interactive endpoint** – An *interactive endpoint* is an HTTPS endpoint to which EMR Studio users can connect a workspace. You can only access the HTTPS endpoints from your EMR Studio, and you create them in a private subnet of the Amazon Virtual Private Cloud (Amazon VPC) for your Amazon EKS cluster.

The Python, PySpark, and Spark Scala kernels use the permissions defined in your Amazon EMR on EKS job execution role to invoke other AWS services. All kernels and users that connect to the interactive endpoint utilize the role that you specified when you created the endpoint. We recommend that you create separate endpoints for different users, and that the users have different AWS Identity and Access Management (IAM) roles.

- **AWS Application Load Balancer controller** – The *AWS Application Load Balancer controller* manages Elastic Load Balancing for an Amazon EKS Kubernetes cluster. The controller provisions an Application Load Balancer (ALB) when you create a Kubernetes Ingress resource. An ALB exposes a Kubernetes service, such as an interactive endpoint, outside of the Amazon EKS cluster but within the same Amazon VPC. When you create an interactive endpoint, an Ingress resource is also deployed that exposes the interactive endpoint by means of the ALB for interactive clients to connect to. You only need to install one AWS Application Load Balancer controller for each Amazon EKS cluster.

The following diagram depicts the interactive endpoints architecture in Amazon EMR on EKS. An Amazon EKS cluster comprises the *compute* to run the analytic workloads, and the *interactive*

endpoint. The Application Load Balancer controller runs in the kube-system namespace; the workloads and interactive endpoints run in the namespace that you specify when you create the virtual cluster. When you create an interactive endpoint, the Amazon EMR on EKS control plane creates the interactive endpoint deployment in the Amazon EKS cluster. Additionally, an instance of the application load balancer ingress is created by the AWS load balancer controller. The application load balancer provides the external interface for clients like EMR Studio to connect to the Amazon EMR cluster and run interactive workloads.



Prerequisites to create an interactive endpoint on Amazon EMR on EKS

This section describes prerequisites to set up an interactive endpoint that EMR Studio can use to connect to an Amazon EMR on EKS cluster and run interactive workloads.

AWS CLI

Follow the steps in [Install or update to the latest version of the AWS CLI](#) to install the latest version of the AWS Command Line Interface (AWS CLI).

Installing eksctl

Follow the steps in [Install kubectl](#) to install the latest version of eksctl. If you are using Kubernetes version 1.22 or later for your Amazon EKS cluster, use an eksctl version greater than 0.117.0.

Amazon EKS cluster

Create an Amazon EKS cluster. Register the cluster as a virtual cluster with Amazon EMR on EKS. The following are requirements and considerations for this cluster:

- The cluster must be in the same Amazon Virtual Private Cloud (VPC) as your EMR Studio.
- The cluster must have at least one private subnet to activate interactive endpoints, to link Git-based repositories, and to launch the Application Load Balancer in private mode.
- There must be at least one private subnet in common between your EMR Studio and the Amazon EKS cluster that you use to register your virtual cluster. This ensures that your interactive endpoint appears as an option in your Studio workspaces, and activates connectivity from Studio to the Application Load Balancer.

There are two methods that you can choose from to connect your Studio and your Amazon EKS cluster:

- Create an Amazon EKS cluster and associate it with the subnets that belong to your EMR Studio.
- Alternatively, create an EMR Studio and specify the private subnets for your Amazon EKS cluster.
- Amazon EKS optimized ARM Amazon Linux AMIs are not supported for Amazon EMR on EKS interactive endpoints.
- Only [Amazon EKS managed node groups](#) and Karpenter provisioned nodes are supported.

Grant Cluster access for Amazon EMR on EKS

Use the steps in [Grant Cluster Access for Amazon EMR on EKS](#) to grant Amazon EMR on EKS access to a specific namespace in your cluster.

Activate IRSA on the Amazon EKS cluster

To activate IAM roles for Service Accounts (IRSA) on the Amazon EKS cluster, follow the steps in [Enable IAM Roles for Service Accounts \(IRSA\)](#).

Create IAM job execution role

You must create an IAM role to run workloads on Amazon EMR on EKS interactive endpoints. We refer to this IAM role as the *job execution role* in this documentation. This IAM role gets assigned to both the interactive endpoint container and the actual execution containers that are created when you submit jobs with EMR Studio. You'll need the Amazon Resource Name (ARN) of your job execution role for Amazon EMR on EKS. There are two steps required for this:

- [Create a IAM role for job execution.](#)
- [Update the trust policy of the job execution role.](#)

Grant users access to Amazon EMR on EKS

The IAM entity (user or role) that makes the request to create an interactive endpoint must also have the following Amazon EC2 and `emr-containers` permissions. Follow the steps described in [Grant users access to Amazon EMR on EKS](#) to grant these permissions that allow Amazon EMR on EKS to create, manage, and delete the security groups that limit inbound traffic to the load balancer of your interactive endpoint.

The following `emr-containers` permissions allow the user to perform basic interactive endpoint operations:

```
"ec2:CreateSecurityGroup",
"ec2:DeleteSecurityGroup",
"ec2:AuthorizeSecurityGroupEgress",
"ec2:AuthorizeSecurityGroupIngress",
"ec2:RevokeSecurityGroupEgress",
"ec2:RevokeSecurityGroupIngress"

"emr-containers:CreateManagedEndpoint",
"emr-containers:ListManagedEndpoints",
"emr-containers:DescribeManagedEndpoint",
"emr-containers>DeleteManagedEndpoint"
```

Register the Amazon EKS cluster with Amazon EMR

Set up a virtual cluster and map it to the namespace in the Amazon EKS cluster where you want to run your jobs. For AWS Fargate-only clusters, use the same namespace for both the Amazon EMR on EKS virtual cluster and Fargate profile.

For information on setting up an Amazon EMR on EKS virtual cluster, see [Register the Amazon EKS cluster with Amazon EMR](#).

Deploy AWS Load Balancer Controller to Amazon EKS cluster

An AWS Application Load Balancer is required for your Amazon EKS cluster. You only need to set up one Application Load Balancer controller per Amazon EKS cluster. For information on setting up the AWS Application Load Balancer controller, see [Installing the AWS Load Balancer Controller add-on](#) in the *Amazon EKS User Guide*.

Creating an interactive endpoint for your virtual cluster

This topic describes a couple ways to create an interactive endpoint using the AWS Command Line Interface (AWS CLI) and includes details on available configuration parameters.

Create an interactive endpoint with the `create-managed-endpoint` command

Specify the parameters in the `create-managed-endpoint` command as follows. Amazon EMR on EKS supports creating interactive endpoints with Amazon EMR releases 6.7.0 and higher.

```
aws emr-containers create-managed-endpoint \
--type JUPYTER_ENTERPRISE_GATEWAY \
--virtual-cluster-id 1234567890abcdef0xxxxxxx \
--name example-endpoint-name \
--execution-role-arn arn:aws:iam::444455556666:role/JobExecutionRole \
--release-label emr-6.9.0-latest \
--configuration-overrides '{
  "applicationConfiguration": [{
    "classification": "spark-defaults",
    "properties": {
      "spark.driver.memory": "2G"
    }
  }],
  "monitoringConfiguration": {
    "cloudWatchMonitoringConfiguration": {
      "logGroupName": "log_group_name",
      "logStreamNamePrefix": "log_stream_prefix"
    },
    "persistentAppUI": "ENABLED",
```

```

        "s3MonitoringConfiguration": {
            "logUri": "s3://my_s3_log_location"
        }
    }
}'

```

For more information, see [Parameters for creating an interactive endpoint](#).

Create an interactive endpoint with specified parameters in a JSON file

1. Create a `create-managed-endpoint-request.json` file and specify the required parameters for your endpoint, as shown in the following JSON file:

```

{
  "name": "MY_TEST_ENDPOINT",
  "virtualClusterId": "MY_CLUSTER_ID",
  "type": "JUPYTER_ENTERPRISE_GATEWAY",
  "releaseLabel": "emr-6.9.0-latest",
  "executionRoleArn": "arn:aws:iam::444455556666:role/JobExecutionRole",
  "configurationOverrides": {
    "applicationConfiguration": [
      {
        "classification": "spark-defaults",
        "properties": {
          "spark.driver.memory": "8G"
        }
      }
    ],
    "monitoringConfiguration": {
      "persistentAppUI": "ENABLED",
      "cloudWatchMonitoringConfiguration": {
        "logGroupName": "my_log_group",
        "logStreamNamePrefix": "log_stream_prefix"
      },
      "s3MonitoringConfiguration": {
        "logUri": "s3://my_s3_log_location"
      }
    }
  }
}

```

```
    }  
  }  
}
```

2. Use the `create-managed-endpoint` command with a path to the `create-managed-endpoint-request.json` file that is stored locally or in Amazon S3.

```
aws emr-containers create-managed-endpoint \  
--cli-input-json file:///./create-managed-endpoint-request.json --region AWS-Region
```

Output of create interactive endpoint

You should see the following output in the terminal. The output includes the name and identifier of your new interactive endpoint:

```
{  
  "id": "1234567890abcdef0",  
  "name": "example-endpoint-name",  
  "arn": "arn:aws:emr-containers:us-west-2:111122223333:/  
virtualclusters/444455556666/endpoints/444455556666",  
  "virtualClusterId": "111122223333xxxxxxxx"  
}
```

Running `aws emr-containers create-managed-endpoint` creates a self-signed certificate that allows HTTPS communication between EMR Studio and the interactive endpoint server.

If you run `create-managed-endpoint` and haven't completed the prerequisites, Amazon EMR returns an error message with the actions that you must take to continue.

Parameters for creating an interactive endpoint

Topics

- [Required parameters for interactive endpoints](#)
- [Optional parameters for interactive endpoints](#)

Required parameters for interactive endpoints

You must specify the following parameters when you create an interactive endpoint:

--type

Use JUPYTER_ENTERPRISE_GATEWAY. This is the only supported type.

--virtual-cluster-id

The identifier of the virtual cluster that you registered with Amazon EMR on EKS.

--name

A descriptive name for the interactive endpoint that helps EMR Studio users select it from the dropdown list.

--execution-role-arn

The Amazon Resource Name (ARN) of your IAM job execution role for Amazon EMR on EKS that was created as part of the prerequisites.

--release-label

The release label of the Amazon EMR release to use for the endpoint. For example, `emr-6.9.0-latest`. Amazon EMR on EKS supports interactive endpoints with Amazon EMR releases 6.7.0 and higher.

Optional parameters for interactive endpoints

Optionally, you can also specify the following parameters when you create an interactive endpoint:

--configuration-overrides

To override the default configurations for applications, supply a configuration object. You can use a shorthand syntax to provide the configuration, or you can reference the configuration object in a JSON file.

Configuration objects consist of a classification, properties, and optional nested configurations. Properties consist of the settings that you want to override in that file. You can specify multiple classifications for multiple applications in a single JSON object. The configuration classifications that are available vary by Amazon EMR on EKS release. For a list of configuration classifications that are available for each release of Amazon EMR on EKS, see [Amazon EMR on EKS releases](#). In addition to the configuration classifications listed for each release, interactive endpoints bring in the additional classification `jeg-config`. For more information, see [Jupyter Enterprise Gateway \(JEG\) configuration options](#).

Configuring settings for interactive endpoints

This section contains a series of topics that cover various configurations for interactive endpoints and pod settings. These give you the ability to monitor and troubleshoot failures, send log information to Amazon S3 or to Amazon CloudWatch Logs, or to create interactive endpoints where you specify custom pod templates.

Topics

- [Monitoring Spark jobs](#)
- [Specifying custom pod templates with interactive endpoints](#)
- [Deploying a JEG pod to a node group](#)
- [Jupyter Enterprise Gateway \(JEG\) configuration options](#)
- [Modifying PySpark session parameters](#)
- [Custom kernel image with interactive endpoint](#)

Monitoring Spark jobs

So that you can monitor and troubleshoot failures, configure your interactive endpoints so that the jobs initiated with the endpoint can send log information to Amazon S3, Amazon CloudWatch Logs, or both. The following sections describe how to send Spark application logs to Amazon S3 for the Spark jobs that you launch with Amazon EMR on EKS interactive endpoints.

Configure IAM policy for Amazon S3 logs

Before your kernels can send log data to Amazon S3, the permissions policy for the job execution role must include the following permissions. Replace *amzn-s3-demo-destination-bucket* with the name of your logging bucket.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
```

```

        "s3:GetObject",
        "s3:ListBucket"
    ],
    "Resource": [
        "arn:aws:s3:::amzn-s3-demo-bucket",
        "arn:aws:s3:::amzn-s3-demo-bucket/*"
    ],
    "Sid": "AllowS3Putobject"
  }
]
}

```

Note

Amazon EMR on EKS can also create an S3 bucket. If an S3 bucket is not available, include the `s3:CreateBucket` permission in the IAM policy.

After you've given your execution role the permissions it needs to send logs to the S3 bucket, your log data is sent to the following Amazon S3 locations. This happens when `s3MonitoringConfiguration` is passed in the `monitoringConfiguration` section of a `create-managed-endpoint` request.

- **Driver logs** – `logUri/virtual-cluster-id/endpoints/endpoint-id/containers/spark-application-id/spark-application-id-driver/(stderr.gz/stdout.gz)`
- **Executor logs** – `logUri/virtual-cluster-id/endpoints/endpoint-id/containers/spark-application-id/executor-pod-name-exec-<Number>/(stderr.gz/stdout.gz)`

Note

Amazon EMR on EKS doesn't upload the endpoint logs to your S3 bucket.

Specifying custom pod templates with interactive endpoints

You can create interactive endpoints where you specify custom pod templates for drivers and executors. *Pod templates* are specifications that determine how to run each pod. You can use pod

template files to define the configurations of driver or executor pods that Spark configurations don't support. Pod templates are currently supported in Amazon EMR releases 6.3.0 and greater.

For more information about pod templates, see [Using pod templates](#) in the *Amazon EMR on EKS Development Guide*.

The following example shows how to create an interactive endpoint with pod templates:

```
aws emr-containers create-managed-endpoint \
  --type JUPYTER_ENTERPRISE_GATEWAY \
  --virtual-cluster-id virtual-cluster-id \
  --name example-endpoint-name \
  --execution-role-arn arn:aws:iam::aws-account-id:role/EKSClusterRole \
  --release-label emr-6.9.0-latest \
  --configuration-overrides '{
    "applicationConfiguration": [
      {
        "classification": "spark-defaults",
        "properties": {
          "spark.kubernetes.driver.podTemplateFile": "path/to/driver/
template.yaml",
          "spark.kubernetes.executor.podTemplateFile": "path/to/executor/
template.yaml"
        }
      }
    ]
  }'
```

Deploying a JEG pod to a node group

JEG (Jupyter Enterprise Gateway) pod placement is a feature that allows you to deploy an interactive endpoint on a specific node group. With this feature, you can configure settings such as instance type for the interactive endpoint.

Associating a JEG pod to a managed node group

The following configuration property allows you to specify the name of a managed node group on your Amazon EKS cluster where the JEG pod will be deployed.

```
//payload
--configuration-overrides '{
  "applicationConfiguration": [
```

```
{
  "classification": "endpoint-configuration",
  "properties": {
    "managed-nodegroup-name": NodeGroupName
  }
}
```

A node group must have the Kubernetes label `for-use-with-emr-containers-managed-endpoint-ng=NodeGroupName` attached to all nodes that are part of the node group. To list all nodes of a node group that have this tag, use the following command:

```
kubectl get nodes --show-labels | grep for-use-with-emr-containers-managed-endpoint-ng=NodeGroupName
```

If the output of the command above doesn't return nodes that are part of your managed node group, then there are no nodes in the node group that have the `for-use-with-emr-containers-managed-endpoint-ng=NodeGroupName` Kubernetes label attached. In this case, follow the steps below to attach that label to the nodes in your node group.

1. Use the following command to add the `for-use-with-emr-containers-managed-endpoint-ng=NodeGroupName` Kubernetes label to all nodes in a managed node group *NodeGroupName*:

```
kubectl label nodes --selector eks:nodegroup-name=NodeGroupName for-use-with-emr-containers-managed-endpoint-ng=NodeGroupName
```

2. Verify that the nodes were labeled correctly using the following command:

```
kubectl get nodes --show-labels | grep for-use-with-emr-containers-managed-endpoint-ng=NodeGroupName
```

A managed node group must be associated with an Amazon EKS cluster's security group, which is usually the case if you created your cluster and managed node group using `eksctl`. You can verify this in the AWS console using the following steps.

1. Go to your cluster in the Amazon EKS console.
2. Go to the networking tab of your cluster and note down the cluster security group.

3. Go to the compute tab of your cluster and click on the managed node group name.
4. Under the **Details** tab of the managed node group, verify that the cluster security group that you noted previously is listed under **Security groups**.

If the managed node group is not attached to the Amazon EKS cluster security group, you need to attach the `for-use-with-emr-containers-managed-endpoint-sg=ClusterName/NodeGroupName` tag to the node group security group. Use the steps below to attach this tag.

1. Go to the Amazon EC2 console and click on security groups on the left navigation pane.
2. Select your managed node group's security group by clicking the checkbox.
3. Under the **Tags** tab, add the tag `for-use-with-emr-containers-managed-endpoint-sg=ClusterName/NodeGroupName` using the **Manage tags** button.

Associating a JEG pod to a self-managed node group

The following configuration property allows you to specify the name of a self-managed or unmanaged node group on the Amazon EKS cluster where the JEG pod will be deployed.

```
//payload
--configuration-overrides '{
  "applicationConfiguration": [
    {
      "classification": "endpoint-configuration",
      "properties": {
        "self-managed-nodegroup-name": NodeGroupName
      }
    }
  ]
}'
```

The node group must have `for-use-with-emr-containers-managed-endpoint-ng=NodeGroupName` Kubernetes label attached to all nodes that are part of the node group. To list all the nodes of a node group that have this tag, use the following command:

```
kubectl get nodes --show-labels | grep for-use-with-emr-containers-managed-endpoint-ng=NodeGroupName
```

If the output of the command above doesn't return nodes that are part of your self-managed node group, then there are no nodes in the node group that have the `for-use-with-emr-containers-managed-endpoint-ng=NodeGroupName` Kubernetes label attached. In this case, follow the steps below to attach that label to the nodes in your node group.

1. If you created the self-managed node group using `eksctl`, then use the following command to add the `for-use-with-emr-containers-managed-endpoint-ng=NodeGroupName` Kubernetes label to all nodes in the self-managed node group `NodeGroupName` at once.

```
kubectl label nodes --selector alpha.eksctl.io/nodegroup-name=NodeGroupName for-use-with-emr-containers-managed-endpoint-ng=NodeGroupName
```

If you didn't use `eksctl` to create the self-managed node group, then you will need to replace the selector in the above command to a different Kubernetes label that is attached to all the nodes of the node group.

2. Use the following command to verify that the nodes were labeled correctly:

```
kubectl get nodes --show-labels | grep for-use-with-emr-containers-managed-endpoint-ng=NodeGroupName
```

The security group for the self-managed node group must have the `for-use-with-emr-containers-managed-endpoint-sg=ClusterName/NodeGroupName` tag attached. Use the following steps to attach the tag to the security group from the AWS Management Console.

1. Navigate to the Amazon EC2 console. Select **Security groups** on the left navigation pane.
2. Select the checkbox next to the security group for your self-managed node group.
3. Under the **Tags** tab, use the **Manage tags** button to add the tag `for-use-with-emr-containers-managed-endpoint-sg=ClusterName/NodeGroupName`. Replace `ClusterName` and `NodeGroupName` with appropriate values.

Associating a JEG pod to a managed node group with On-Demand instances

You can also define additional labels, known as *Kubernetes label selectors*, to specify additional constraints or restrictions to run an interactive endpoint on a given node or node group. The following example shows how to use On-Demand Amazon EC2 instances for a JEG pod.

```
--configuration-overrides '{
  "applicationConfiguration": [
    {
      "classification": "endpoint-configuration",
      "properties": {
        "managed-nodegroup-name": NodeGroupName,
        "node-labels": "eks.amazonaws.com/capacityType:ON_DEMAND"
      }
    }
  ]
}'
```

Note

You can only use the `node-labels` property with either with a `managed-nodegroup-name` or `self-managed-nodegroup-name` property.

Jupyter Enterprise Gateway (JEG) configuration options

Amazon EMR on EKS uses Jupyter Enterprise Gateway (JEG) to turn on interactive endpoints. You can set the following values for the allow-listed JEG configurations when you create the endpoint.

- **`RemoteMappingKernelManager.cull_idle_timeout`** – Timeout in seconds (integer), after which a kernel is considered idle and ready to be culled. Values of 0 or lower deactivate culling. Short timeouts might result in kernels being culled for users with poor network connections.
- **`RemoteMappingKernelManager.cull_interval`** – The interval in seconds (integer) on which to check for idle kernels that exceed the cull timeout value.

Modifying PySpark session parameters

Starting with Amazon EMR on EKS release 6.9.0, in Amazon EMR Studio you can adjust the Spark configuration associated with a PySpark session by executing the `%%configure` magic command in the EMR notebook cell.

The following example shows a sample payload that you can use to modify memory, cores, and other properties for the Spark driver and executor. For the `conf` settings, you can configure any Spark configuration mentioned in the [Apache Spark configuration documentation](#).


```
%%configure -f
{
  "driverMemory": "16G",
  "driverCores": 4,
  "executorMemory" : "32G",
  "executorCores": 2,
  "conf": {
    "spark.dynamicAllocation.maxExecutors" : 10,
    "spark.dynamicAllocation.minExecutors": 1
  }
}
```

The following example shows a sample payload that you can use to add files, pyFiles, and jar dependencies to a Spark runtime.

```
%%configure -f
{
  "files": "s3://amzn-s3-demo-bucket-emr-eks/sample_file.txt",
  "pyFiles": : "path-to-python-files",
  "jars" : "path-to-jars"
}
```

Custom kernel image with interactive endpoint

To ensure that you have the correct dependencies for your application when you run interactive workloads from Amazon EMR Studio, you can customize Docker images for interactive endpoints and run customized base kernel images. To create an interactive endpoint and connect it with a custom Docker image, perform the following steps.

Note

You can only override base images. You can't add new kernel image types.

1. **Create and publish a customized Docker image.** The base image contains the Spark runtime and the notebook kernels that run with it. To create the image, you can follow steps 1 through 4 in [How to customize Docker images](#). In step 1, the base image URI in your Docker file must use `notebook-spark` in place of `spark`.

```
ECR-registry-account.dkr.ecr.Region.amazonaws.com/notebook-spark/container-image-tag
```

For more information on how to select AWS Regions and container image tags, see [Details for selecting a base image URL](#).

2. Create an interactive endpoint that can be used with the custom image.

- a. Create a JSON file `custom-image-managed-endpoint.json` with the following contents. This example uses Amazon EMR release 6.9.0.

Example

```
{
  "name": "endpoint-name",
  "virtualClusterId": "virtual-cluster-id",
  "type": "JUPYTER_ENTERPRISE_GATEWAY",
  "releaseLabel": "emr-6.9.0-latest",
  "executionRoleArn": "execution-role-arn",
  "configurationOverrides": {
    "applicationConfiguration": [
      {
        "classification": "jupyter-kernel-overrides",
        "configurations": [
          {
            "classification": "python3",
            "properties": {
              "container-image": "123456789012.dkr.ecr.us-west-2.amazonaws.com/custom-notebook-python:latest"
            }
          },
          {
            "classification": "spark-python-kubernetes",
            "properties": {
              "container-image": "123456789012.dkr.ecr.us-west-2.amazonaws.com/custom-notebook-spark:latest"
            }
          }
        ]
      }
    ]
  }
}
```

```
}
```

- b. Create an interactive endpoint with the configurations specified in the JSON file as shown in the following example. For more information, see [Create an interactive endpoint with the create-managed-endpoint command](#).

```
aws emr-containers create-managed-endpoint --cli-input-json custom-image-managed-endpoint.json
```

3. **Connect to the interactive endpoint via EMR Studio.** For more information and steps to complete, see [Connecting from Studio](#) in the Amazon EMR on EKS section of the AWS Workshop Studio docs.

Monitoring interactive endpoints

With Amazon EMR on EKS version 6.10 and later, interactive endpoints emit Amazon CloudWatch metrics for monitoring and troubleshooting kernel lifecycle operations. Metrics are triggered by interactive clients, such as EMR Studio or self-hosted Jupyter notebooks. Each of the operations supported by interactive endpoints have metrics associated with them. The operations are modeled as dimensions to each metric, as shown in the table below. Metrics emitted by interactive endpoints are visible under a custom namespace, EMRContainers, in your account.

Metric	Description	Unit
RequestCount	Cumulative number of requests of an operation processed by the interactive endpoint.	Count
RequestLatency	The time from when a request arrived at the interactive endpoint and a response was sent by the interactive endpoint.	Millisecond
4XXError	Emitted when a request for an operation results in a 4xx error during processing.	Count

Metric	Description	Unit
5XXError	Emitted when a request for an operation results in a 5Xxx server side error.	Count
KernelLaunchSuccess	Applicable only for the CreateKernel operation. It indicates the cumulative number of kernel launches that were successful up to and including this request.	Count
KernelLaunchFailure	Applicable only for the CreateKernel operation. It indicates the cumulative number of kernel launch failures up until and including this request.	Count

Each interactive endpoint metric has the following dimensions attached to it:

- **ManagedEndpointId** – Identifier for the interactive endpoint
- **OperationName** – The operation triggered by the interactive client

Possible values for the **OperationName** dimension are shown in the following table:

operationName	Operation description
CreateKernel	Request that the interactive endpoint start a kernel.
ListKernels	Request that the interactive endpoint list the kernels that have been previously started using the same session token.

operationName	Operation description
GetKernel	Request that the interactive endpoint get details about a specific kernel that has been previously started.
ConnectKernel	Request that the interactive endpoint establish connectivity between the notebook client and the kernel.
ConfigureKernel	Publish %%configure magic request on a pyspark kernel.
ListKernelSpecs	Request that the interactive endpoint list the available kernel specs.
GetKernelSpec	Request that the interactive endpoint get the kernel specs of a kernel that has been previously launched.
GetKernelSpecResource	Request that the interactive endpoint get specific resources associated with the kernel specs that have been previously launched.

Examples

To access the total number of kernels launched for an interactive endpoint on a given day:

1. Select the custom namespace: EMRContainers
2. Select your ManagedEndpointId, OperationName – CreateKernel
3. RequestCount metric with the statistic SUM and period 1 day will provide all the kernel launch requests made in the last 24 hours.
4. KernelLaunchSuccess metric with statistic SUM and period 1 day will provide all the successful kernel launch requests made in the last 24 hours.

To access the number of kernel failures for an interactive endpoint on a given day:

1. Select the custom namespace: EMRContainers
2. Select your ManagedEndpointId, OperationName - CreateKernel
3. KernelLaunchFailure metric with statistic SUM and period 1 day will provide all the failed kernel launch requests made in the last 24 hours. You can also select the 4XXError and 5XXError metric to know what kind of kernel launch failure happened.

Using self-hosted Jupyter notebooks

You can host and manage Jupyter or JupyterLab notebooks on an Amazon EC2 instance or on your own Amazon EKS cluster as a *self-hosted Jupyter notebook*. You can then run interactive workloads with your self-hosted Jupyter notebooks. The following sections walk through the process to set up and deploy a self-hosted Jupyter notebook on an Amazon EKS cluster.

Creating a self-hosted Jupyter notebook on an EKS cluster

- [Create a security group](#)
- [Create an Amazon EMR on EKS interactive endpoint](#)
- [Retrieve the gateway server URL of your interactive endpoint](#)
- [Retrieve an auth token to connect to the interactive endpoint](#)
- [Example: Deploy a JupyterLab notebook](#)
- [Delete a self-hosted Jupyter notebook](#)

Create a security group

Before you can create an interactive endpoint and run a self-hosted Jupyter or JupyterLab notebook, you must create a security group to control the traffic between your notebook and the interactive endpoint. To use the Amazon EC2 console or Amazon EC2 SDK to create the security group, refer to the steps in [Create a security group](#) in the *Amazon EC2 User Guide*. You should create the security group in the VPC where you want to deploy your notebook server.

To follow the example in this guide, use the same VPC as your Amazon EKS cluster. If you want to host your notebook in a VPC that is different from the VPC for your Amazon EKS cluster, you

might need to create a peering connection between those two VPCs. For steps to create a peering connection between two VPCs, see [Create a VPC peering connection](#) in the Amazon VPC Getting Started Guide.

You need the ID for the security group to [create an Amazon EMR on EKS interactive endpoint](#) in the next step.

Create an Amazon EMR on EKS interactive endpoint

After you create security group for your notebook, use the steps provided in [Creating an interactive endpoint for your virtual cluster](#) to create an interactive endpoint. You must provide the security group ID that you created for your notebook in [Create a security group](#).

Insert the security ID in place of *your-notebook-security-group-id* in the following configuration override settings:

```
--configuration-overrides '{
  "applicationConfiguration": [
    {
      "classification": "endpoint-configuration",
      "properties": {
        "notebook-security-group-id": "your-notebook-security-group-id"
      }
    }
  ],
  "monitoringConfiguration": {
    ...'
```

Retrieve the gateway server URL of your interactive endpoint

After you create an interactive endpoint, retrieve the gateway server URL with the `describe-managed-endpoint` command in the AWS CLI. You need this URL to connect your notebook to the endpoint. The gateway server URL is a private endpoint.

```
aws emr-containers describe-managed-endpoint \
--region region \
--virtual-cluster-id virtualClusterId \
--id endpointId
```

Initially, your endpoint is in the **CREATING** state. After a few minutes, it transitions to the **ACTIVE** state. When the endpoint is **ACTIVE**, it's ready to use.

Take note of the `serverUrl` attribute that the `aws emr-containers describe-managed-endpoint` command returns from the active endpoint. You need this URL to connect your notebook to the endpoint when you [deploy your self-hosted Jupyter or JupyterLab notebook](#).

Retrieve an auth token to connect to the interactive endpoint

To connect to an interactive endpoint from a Jupyter or JupyterLab notebook, you must generate a session token with the `GetManagedEndpointSessionCredentials` API. The token acts as proof of authentication to connect to the interactive endpoint server.

The following command is explained in more detail with an output example below.

```
aws emr-containers get-managed-endpoint-session-credentials \
--endpoint-identifier endpointArn \
--virtual-cluster-identifier virtualClusterArn \
--execution-role-arn executionRoleArn \
--credential-type "TOKEN" \
--duration-in-seconds durationInSeconds \
--region region
```

endpointArn

The ARN of your endpoint. You can find the ARN in the result of a `describe-managed-endpoint` call.

virtualClusterArn

The ARN of the virtual cluster.

executionRoleArn

The ARN of the execution role.

durationInSeconds

The duration in seconds for which the token is valid. The default duration is 15 minutes (900), and the maximum is 12 hours (43200).

region

The same region as your endpoint.

Your output should resemble the following example. Take note of the *session-token* value that you will use when you [deploy your self-hosted Jupyter or JupyterLab notebook](#).


```
{
  "id": "credentialsId",
  "credentials": {
    "token": "session-token"
  },
  "expiresAt": "2022-07-05T17:49:38Z"
}
```

Example: Deploy a JupyterLab notebook

Once you've completed the steps above, you can try this example procedure to deploy a JupyterLab notebook into the Amazon EKS cluster with your interactive endpoint.

1. Create a namespace to run the notebook server.
2. Create a file locally, `notebook.yaml`, with the following contents. The file contents are described below.

```
apiVersion: v1
kind: Pod
metadata:
  name: jupyter-notebook
  namespace: namespace
spec:
  containers:
    - name: minimal-notebook
      image: jupyter/all-spark-notebook:lab-3.1.4 # open source image
      ports:
        - containerPort: 8888
      command: ["start-notebook.sh"]
      args: ["--LabApp.token=''"]
      env:
        - name: JUPYTER_ENABLE_LAB
          value: "yes"
        - name: KERNEL_LAUNCH_TIMEOUT
          value: "400"
        - name: JUPYTER_GATEWAY_URL
          value: "serverUrl"
        - name: JUPYTER_GATEWAY_VALIDATE_CERT
          value: "false"
        - name: JUPYTER_GATEWAY_AUTH_TOKEN
          value: "session-token"
```

If you are deploying Jupyter notebook to a Fargate-only cluster, label the Jupyter pod with a `role` label as shown in the following example:

```
...
metadata:
  name: jupyter-notebook
  namespace: default
  labels:
    role: example-role-name-label
spec:
  ...
```

namespace

The Kubernetes namespace that the notebook deploys into.

serverUrl

The `serverUrl` attribute that the `describe-managed-endpoint` command returned in [Retrieve the gateway server URL of your interactive endpoint](#).

session-token

The `session-token` attribute that the `get-managed-endpoint-session-credentials` command returned in [Retrieve an auth token to connect to the interactive endpoint](#).

KERNEL_LAUNCH_TIMEOUT

The amount of time in seconds that the interactive endpoint waits for the kernel to come to **RUNNING** state. Ensure sufficient time for kernel launch to complete by setting the kernel launch timeout to an appropriate value (maximum 400 seconds).

KERNEL_EXTRA_SPARK_OPTS

Optionally, you can pass additional Spark configurations for the Spark kernels. Set this environment variable with the values as the Spark configuration property as shown in the following example:

```
- name: KERNEL_EXTRA_SPARK_OPTS
  value: "--conf spark.driver.cores=2
         --conf spark.driver.memory=2G
```

```
--conf spark.executor.instances=2
--conf spark.executor.cores=2
--conf spark.executor.memory=2G
--conf spark.dynamicAllocation.enabled=true
--conf spark.dynamicAllocation.shuffleTracking.enabled=true
--conf spark.dynamicAllocation.minExecutors=1
--conf spark.dynamicAllocation.maxExecutors=5
--conf spark.dynamicAllocation.initialExecutors=1
"
```

3. Deploy the pod spec to your Amazon EKS cluster:

```
kubectl apply -f notebook.yaml -n namespace
```

This will start up a minimal JupyterLab notebook connected to your Amazon EMR on EKS interactive endpoint. Wait until the pod is **RUNNING**. You can check its status with the following command:

```
kubectl get pod jupyter-notebook -n namespace
```

When the pod is ready, the `get pod` command returns output similar to this:

NAME	READY	STATUS	RESTARTS	AGE
jupyter-notebook	1/1	Running	0	46s

4. Attach the notebook security group to the node where the notebook is scheduled.

- a. First, identify the node where `jupyter-notebook` pod is scheduled with the `describe pod` command.

```
kubectl describe pod jupyter-notebook -n namespace
```


- b. Open the Amazon EKS console at <https://console.aws.amazon.com/eks/home#/clusters>.
- c. Navigate to the **Compute** tab for your Amazon EKS cluster and select the node identified by the `describe pod` command. Select the instance ID for the node.
- d. From the **Actions** menu, select **Security** > **Change security groups** to attach the security group that you created in [Create a security group](#).
- e. If you are deploying Jupyter notebook pod on AWS Fargate, create a [SecurityGroupPolicy](#) to apply to the Jupyter notebook pod with the role label:

```
cat >my-security-group-policy.yaml <<EOF
apiVersion: vpcresources.k8s.aws/v1beta1
kind: SecurityGroupPolicy
metadata:
  name: example-security-group-policy-name
  namespace: default
spec:
  podSelector:
    matchLabels:
      role: example-role-name-label
  securityGroups:
    groupIds:
      - your-notebook-security-group-id
EOF
```

5. Now, port-forward so that you can locally access the JupyterLab interface:

```
kubectl port-forward jupyter-notebook 8888:8888 -n namespace
```

Once that is running, navigate to your local browser and visit `localhost:8888` to see the JupyterLab interface:



6. From JupyterLab, create a new Scala notebook. Here is a sample code snippet that you can run to approximate the value of Pi:

```
import scala.math.random
import org.apache.spark.sql.SparkSession

/** Computes an approximation to pi */
val session = SparkSession
  .builder
  .appName("Spark Pi")
  .getOrCreate()

val slices = 2
// avoid overflow
val n = math.min(100000L * slices, Int.MaxValue).toInt


val count = session.sparkContext
  .parallelize(1 until n, slices)
  .map { i =>
    val x = random * 2 - 1
    val y = random * 2 - 1
```

```

    if (x*x + y*y <= 1) 1 else 0
  }.reduce(_ + _)

println(s"Pi is roughly ${4.0 * count / (n - 1)}")
session.stop()

```



Delete a self-hosted Jupyter notebook

When you're ready to delete your self-hosted notebook, you can also delete the interactive endpoint and security group, too. Perform the actions in the following order:

1. Use the following command to delete the `jupyter-notebook` pod:

```
kubectl delete pod jupyter-notebook -n namespace
```

2. Then, delete your interactive endpoint with the `delete-managed-endpoint` command. For steps to delete an interactive endpoint, see [Delete an interactive endpoint](#). Initially, your endpoint will be in the **TERMINATING** state. Once all resources have been cleaned up, it transitions to the **TERMINATED** state.

3. If you don't plan to use the notebook security group that you created in [Create a security group](#) for other Jupyter notebook deployments, you can delete it. See [Delete a security group](#) in the Amazon EC2 User Guide for more information.

Getting information about interactive endpoints with CLI commands

This topic covers the supported operations on an interactive endpoint other than [create-managed-endpoint](#).

Fetch interactive endpoint details

After you create an interactive endpoint, you can retrieve its details using the `describe-managed-endpoint` AWS CLI command. Insert your own values for *managed-endpoint-id*, *virtual-cluster-id*, and *region*:

```
aws emr-containers describe-managed-endpoint --id managed-endpoint-id \
--virtual-cluster-id virtual-cluster-id --region region
```

The output looks similar to the following, with the specified endpoint, such as ARN, ID, and name.

```
{
  "id": "as3ys2xxxxxxx",
  "name": "endpoint-name",
  "arn": "arn:aws:emr-containers:us-east-1:1828xxxxxxx:/virtualclusters/
lbhl6kwwyxxxxxxxxxxxxxxxxx/endpoints/as3ysxxxxxxxx",
  "virtualClusterId": "lbhl6kwwyxxxxxxxxxxxxxxxxx",
  "type": "JUPYTER_ENTERPRISE_GATEWAY",
  "state": "ACTIVE",
  "releaseLabel": "emr-6.9.0-latest",
  "executionRoleArn": "arn:aws:iam::1828xxxxxxx:role/RoleName",
  "certificateAuthority": {
    "certificateArn": "arn:aws:acm:us-east-1:1828xxxxxxx:certificate/zzzzzzzz-
e59b-4ed0-aaaa-bbbbbbbbbbbbbb",
    "certificateData": "certificate-data"
  },
  "configurationOverrides": {
    "applicationConfiguration": [
      {
        "classification": "spark-defaults",
```

```

        "properties": {
            "spark.driver.memory": "8G"
        }
    ],
    "monitoringConfiguration": {
        "persistentAppUI": "ENABLED",
        "cloudWatchMonitoringConfiguration": {
            "logGroupName": "log-group-name",
            "logStreamNamePrefix": "log-stream-name-prefix"
        },
        "s3MonitoringConfiguration": {
            "logUri": "s3-bucket-name"
        }
    },
    "serverUrl": "https://internal-k8s-namespace-ingressa-aaaaaaaaa-
zzzzzzzzzz.us-east-1.elb.amazonaws.com:18888 (https://internal-k8s-nspluto-
ingressa-51e860abbd-1620715833.us-east-1.elb.amazonaws.com:18888/)",
    "createdAt": "2022-09-19T12:37:49+00:00",
    "securityGroup": "sg-aaaaaaaaaaaaaa",
    "subnetIds": [
        "subnet-111111111111",
        "subnet-222222222222",
        "subnet-333333333333"
    ],
    "stateDetails": "Endpoint created successfully. It took 3 Minutes 15 Seconds",
    "tags": {}
}

```

List all interactive endpoints associated with a virtual cluster

Use the `list-managed-endpoints` AWS CLI command to fetch a list of all the interactive endpoints associated with a specified virtual cluster. Replace `virtual-cluster-id` with the ID of your virtual cluster.

```
aws emr-containers list-managed-endpoints --virtual-cluster-id virtual-cluster-id
```

The output of the `list-managed-endpoint` command is shown below:

```
{
  "endpoints": [{
```



```

    "id": "as3ys2xxxxxxx",
    "name": "endpoint-name",
    "arn": "arn:aws:emr-containers:us-east-1:1828xxxxxxx:/virtualclusters/
lbhl6kwwyoxxxxxxxxxxxxxxxxx/endpoints/as3ysxxxxxxxxx",
    "virtualClusterId": "lbhl6kwwyoxxxxxxxxxxxxxxxxx",
    "type": "JUPYTER_ENTERPRISE_GATEWAY",
    "state": "ACTIVE",
    "releaseLabel": "emr-6.9.0-latest",
    "executionRoleArn": "arn:aws:iam::1828xxxxxxx:role/RoleName",
    "certificateAuthority": {
        "certificateArn": "arn:aws:acm:us-east-1:1828xxxxxxx:certificate/zzzzzzzz-
e59b-4ed0-aaaa-bbbbbbbbbbbb",
        "certificateData": "certificate-data"
    },
    "configurationOverrides": {
        "applicationConfiguration": [{
            "classification": "spark-defaults",
            "properties": {
                "spark.driver.memory": "8G"
            }
        }],
        "monitoringConfiguration": {
            "persistentAppUI": "ENABLED",
            "cloudWatchMonitoringConfiguration": {
                "logGroupName": "log-group-name",
                "logStreamNamePrefix": "log-stream-name-prefix"
            },
            "s3MonitoringConfiguration": {
                "logUri": "s3-bucket-name"
            }
        }
    },
    "serverUrl": "https://internal-k8s-namespace-ingressa-aaaaaaaaa-
zzzzzzzzzz.us-east-1.elb.amazonaws.com:18888 (https://internal-k8s-nspluto-
ingressa-51e860abbd-1620715833.us-east-1.elb.amazonaws.com:18888/)",
    "createdAt": "2022-09-19T12:37:49+00:00",
    "securityGroup": "sg-aaaaaaaaaaaaa",
    "subnetIds": [
        "subnet-1111111111",
        "subnet-2222222222",
        "subnet-3333333333"
    ],
    "stateDetails": "Endpoint created successfully. It took 3 Minutes 15 Seconds",
    "tags": {}

```

```
    }]  
}
```

Delete an interactive endpoint

To delete an interactive endpoint associated with an Amazon EMR on EKS virtual cluster, use the `delete-managed-endpoint` AWS CLI command. When you delete an interactive endpoint, Amazon EMR on EKS removes the default security groups that were created for that endpoint.

Specify values for the following parameters to the command:

- **--id:** The identifier of the interactive endpoint that you want to delete.
- **--virtual-cluster-id** – The identifier of the virtual cluster associated with the interactive endpoint that you want to delete. This is the same virtual cluster ID that was specified when the interactive endpoint was created.

```
aws emr-containers delete-managed-endpoint --id managed-endpoint-id --virtual-cluster-id virtual-cluster-id
```

The command returns output similar to the following to confirm that you deleted the interactive endpoint:

```
{  
  "id": "8gai4l4exxxxxx",  
  "virtualClusterId": "0b0qvauoy3ch1nqodxxxxxxxx"  
}
```

Read, Write and Upload data into Amazon S3 Express One Zone with Amazon EMR on EKS

With Amazon EMR releases 7.2.0 and higher, you can use Amazon EMR on EKS with the [Amazon S3 Express One Zone](#) storage class for improved performance when you run jobs and workloads. S3 Express One Zone is a high-performance, single-zone Amazon S3 storage class that delivers consistent, single-digit millisecond data access for most latency-sensitive applications. At the time of its release, S3 Express One Zone delivers the lowest latency and highest performance cloud object storage in Amazon S3.

Prerequisites

Before you can use S3 Express One Zone with Amazon EMR on EKS, you must have the following prerequisites:

- [Completed setting up Amazon EMR on EKS](#).
- After you set up Amazon EMR on EKS, [create a virtual cluster](#).

Getting started with S3 Express One Zone

Follow these steps to get started with S3 Express One Zone

1. Add the `CreateSession` permission to your job execution role. When S3 Express One Zone initially performs an action like `GET`, `LIST`, or `PUT` on an S3 object, the storage class calls `CreateSession` on your behalf. The following is an example of how to grant the `CreateSession` permission.

JSON

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

```

        "arn:aws:s3express:*:*:bucket/DOC-EXAMPLE-BUCKET"
      ],
      "Action": [
        "s3express:CreateSession"
      ],
      "Sid": "AllowS3EXPRESSCreatesession"
    }
  ]
}

```

2. You must use the Apache Hadoop connector S3A to access the S3 Express buckets, so change your Amazon S3 URIs to use the s3a scheme to use the connector. If they don't use the scheme, you can change the filesystem implementation that you use for s3 and s3n schemes.

To change the s3 scheme, specify the following cluster configurations:

```

[
  {
    "Classification": "core-site",
    "Properties": {
      "fs.s3.impl": "org.apache.hadoop.fs.s3a.S3AFileSystem",
      "fs.AbstractFileSystem.s3.impl": "org.apache.hadoop.fs.s3a.S3A"
    }
  }
]

```

To change the s3n scheme, specify the following cluster configurations:

```

[
  {
    "Classification": "core-site",
    "Properties": {
      "fs.s3n.impl": "org.apache.hadoop.fs.s3a.S3AFileSystem",
      "fs.AbstractFileSystem.s3n.impl": "org.apache.hadoop.fs.s3a.S3A",
      "fs.s3a.endpoint.region": "us-west-2",
      "fs.s3a.change.detection.mode": "none",
      "fs.s3a.select.enabled": "false"
    }
  },
  {
    "Classification": "spark-defaults",
    "Properties": {

```

```
"spark.hadoop.fs.s3a.aws.credentials.provider":  
  "software.amazon.awssdk.auth.credentials.WebIdentityTokenFileCredentialsProvider",  
  "spark.sql.sources.fastS3PartitionDiscovery.enabled": "false"  
  }  
}  
]
```

3. In your spark-submit configuration, use the web identity credential provider.

```
"spark.hadoop.fs.s3a.aws.credentials.provider=com.amazonaws.auth.WebIdentityTokenCredential
```

Monitoring jobs

You can use Amazon CloudWatch Events to track jobs that run on an Amazon EMR on EKS virtual cluster. You can use events to track the activity and health of a jobs that you run on a virtual cluster. The topics that follow show you ways to configure monitoring effectively to maintain the health of your resources.

Topics

- [Monitor jobs with Amazon CloudWatch Events](#)
- [Automate Amazon EMR on EKS with CloudWatch Events](#)
- [Example: Set up a rule that invokes Lambda](#)
- [Monitor job's driver pod with a retry policy using Amazon CloudWatch Events](#)

Monitor jobs with Amazon CloudWatch Events

Amazon EMR on EKS emits events when the state of a job run changes. Each event provides information, such as the date and time when the event occurred, along with further details about the event, such as the virtual cluster ID and the ID of the job run that was affected.

You can use events to track the activity and health of a jobs that you run on a virtual cluster. You can also use Amazon CloudWatch Events to define an action to take when a job run generates an event that matches a pattern that you specify. Events are useful for monitoring a specific occurrence during the lifecycle of a job run. For example, you can monitor when a job run changes state from submitted to running. For more information about CloudWatch Events, see the [Amazon EventBridge User Guide](#).

The following table lists Amazon EMR on EKS events along with the state or state change that the event indicates, the severity of the event, and event messages. Each event is represented as a JSON object that is sent automatically to an event stream. The JSON object includes further details about the event. The JSON object is particularly important when you set up rules for event processing using CloudWatch Events because rules seek to match patterns in the JSON object. For more information, see [Amazon EventBridge event patterns](#) and Amazon EMR on EKS Events in the [Amazon EventBridge User Guide](#).

Job run state change events

State	Severity	Message
SUBMITTED	INFO	Job Run <i>JobRunId</i> (<i>JobRunName</i>) was successfully submitted to virtual cluster <i>VirtualClusterId</i> at <i>Time</i> UTC.
RUNNING	INFO	Job Run <i>JobRunId</i> (<i>JobRunName</i>) in virtual cluster <i>VirtualClusterId</i> started running at <i>Time</i> .
COMPLETED	INFO	Job Run <i>jobRunId</i> (<i>JobRunName</i>) in virtual cluster <i>VirtualClusterId</i> completed at <i>Time</i> . The Job Run started running at <i>Time</i> and took <i>Num</i> minutes to complete.
CANCELLED	WARN	Cancellation request has succeeded for Job Run <i>JobRunId</i> (<i>JobRunName</i>) in virtual cluster <i>VirtualClusterId</i> at <i>Time</i> and the Job Run is now cancelled.
FAILED	ERROR	Job Run <i>JobRunId</i> (<i>JobRunName</i>) in virtual cluster <i>VirtualClusterId</i> failed at <i>Time</i> .

Automate Amazon EMR on EKS with CloudWatch Events

You can use Amazon CloudWatch Events to automate your AWS services to respond to system events such as application availability issues or resource changes. Events from AWS services are delivered to CloudWatch Events in near real time. You can write simple rules to indicate which events are of interest to you and what automated actions to take when an event matches a rule. The actions that can be automatically triggered include the following:

- Invoking an AWS Lambda function
- Invoking Amazon EC2 Run Command
- Relaying the event to Amazon Kinesis Data Streams
- Activating an AWS Step Functions state machine

- Notifying an Amazon Simple Notification Service (SNS) topic or an Amazon Simple Queue Service (SQS) queue

Some examples of using CloudWatch Events with Amazon EMR on EKS include the following:

- Activating a Lambda function when a job run succeeds
- Notifying an Amazon SNS topic when a job run fails

CloudWatch Events for "detail-type:" "EMR Job Run State Change" are generated by Amazon EMR on EKS for SUBMITTED, RUNNING, CANCELLED, FAILED and COMPLETED state changes.

Example: Set up a rule that invokes Lambda

Use the following steps to set up a CloudWatch Events rule that invokes Lambda when there is an "EMR Job Run State Change" event.

```
aws events put-rule \  
--name cwe-test \  
--event-pattern '{"detail-type": ["EMR Job Run State Change"]}'
```

Add the Lambda function that you own as a new target and give CloudWatch Events permission to invoke the Lambda function as follows. Replace **123456789012** with your account ID.

```
aws events put-targets \  
--rule cwe-test \  
--targets Id=1,Arn=arn:aws:lambda:us-east-1:123456789012:function:MyFunction
```

```
aws lambda add-permission \  
--function-name MyFunction \  
--statement-id MyId \  
--action 'lambda:InvokeFunction' \  
--principal events.amazonaws.com
```


Note

You cannot write a program that depends on the order or existence of notification events, as they might be out of sequence or missing. Events are emitted on a best effort basis.

Monitor job's driver pod with a retry policy using Amazon CloudWatch Events

Using CloudWatch events, you can monitor driver pods that have been created in jobs that have retry policies. For more information, see [Monitoring a job with a retry policy](#) in this guide.

Managing virtual clusters

A virtual cluster is a Kubernetes namespace that Amazon EMR is registered with. You can create, describe, list, and delete virtual clusters. They do not consume any additional resource in your system. A single virtual cluster maps to a single Kubernetes namespace. Given this relationship, you can model virtual clusters the same way you model Kubernetes namespaces to meet your requirements. See possible use cases in the [Kubernetes Concepts Overview](#) documentation.

To register Amazon EMR with a Kubernetes namespace on an Amazon EKS cluster, you need the name of the EKS cluster and the namespace that has been set up for running your workload. These registered clusters in Amazon EMR are called virtual clusters because they do not manage physical compute or storage but point to a Kubernetes namespace where your workload is scheduled.

Note

Before creating a virtual cluster, you must first complete the steps 1-8 in [Setting up Amazon EMR on EKS](#).

Topics

- [Create a virtual cluster](#)
- [List virtual clusters](#)
- [Describe a virtual cluster](#)
- [Delete a virtual cluster](#)
- [Virtual cluster states](#)
- [Concurrent job limits for virtual clusters](#)

Create a virtual cluster

Run the following command to create a virtual cluster by registering Amazon EMR with a namespace on an EKS cluster. Replace *virtual_cluster_name* with a name that you provide for your virtual cluster. Replace *eks_cluster_name* with the name of the EKS cluster. Replace the *namespace_name* with the namespace that you want to register Amazon EMR with.

```
aws emr-containers create-virtual-cluster \
```

```
--name virtual_cluster_name \  
--container-provider '{  
    "id": "eks_cluster_name",  
    "type": "EKS",  
    "info": {  
        "eksInfo": {  
            "namespace": "namespace_name"  
        }  
    }  
}'
```

Alternatively, you can create a JSON file that includes the required parameters for the virtual cluster, as the following example demonstrates.

```
{  
    "name": "virtual_cluster_name",  
    "containerProvider": {  
        "type": "EKS",  
        "id": "eks_cluster_name",  
        "info": {  
            "eksInfo": {  
                "namespace": "namespace_name"  
            }  
        }  
    }  
}
```

Then run the following `create-virtual-cluster` command with the path to the JSON file.

```
aws emr-containers create-virtual-cluster \  
--cli-input-json file:///./create-virtual-cluster-request.json
```

Note

To validate the successful creation of a virtual cluster, view the status of virtual clusters by running the `list-virtual-clusters` command or by going to the **Virtual clusters** page in the Amazon EMR console.

List virtual clusters

Run the following command to view the status of virtual clusters.

```
aws emr-containers list-virtual-clusters
```

Describe a virtual cluster

Run the following command to get more details about a virtual cluster, such as namespace, status, and date registered. Replace **123456** with your virtual cluster ID.

```
aws emr-containers describe-virtual-cluster --id 123456
```

Delete a virtual cluster

Run the following command to delete a virtual cluster. Replace **123456** with your virtual cluster ID.

```
aws emr-containers delete-virtual-cluster --id 123456
```

Virtual cluster states

The following table describes the four possible states of a virtual cluster.

State	Description
RUNNING	Virtual cluster is in RUNNING state.
TERMINATING	The requested termination of the virtual cluster is in progress.
TERMINATED	The requested termination is complete.
ARRESTED	The requested termination failed because of insufficient permissions.

Concurrent job limits for virtual clusters

You can configure concurrent job limits on an Amazon EMR on EKS virtual cluster to control how many job runs execute simultaneously and how many can wait in queue. You set the concurrency limit (`maxConcurrentJobRuns`) and the queue depth (`maxInQueueJobRuns`) independently, so you can cap running job runs, queued job runs, or both. When you set these limits, the `StartJobRun` API provides backpressure at the virtual cluster level. Job runs beyond the running limit wait in the queue in the `SUBMITTED` or `PENDING` state instead of starting immediately, and once the queue is full, `StartJobRun` rejects further submissions. For example, if you set a virtual cluster to allow 500 concurrent job runs and 100 queued job runs, the 101st queued submission is rejected, and you can rebalance that workload across other virtual clusters on the same EKS cluster or add capacity. When you have not set a concurrency limit and queue depth keeps growing, so that job runs stay in the `SUBMITTED` or `PENDING` state longer before they start, it can signal that the underlying EKS cluster is running low on compute resources and cannot schedule new pods fast enough. In that case, route the workload to another cluster or add capacity.

Concurrent job limits add a control layer in front of the Kubernetes scheduler and the [ResourceQuota](#) feature on the Kubernetes website. Because they are enforced at `StartJobRun`, before any pods are created, excess load is queued or rejected at the API, which protects the underlying cluster before jobs ever reach it. Kubernetes still enforces the actual CPU and memory ceiling underneath.

Key benefits of concurrent job limits

- **Prevents noisy-neighbor overload** — Limits the number of running and queued job runs per virtual cluster, so that a single virtual cluster cannot monopolize the shared EKS cluster and cause noisy-neighbor scheduling failures for other virtual clusters.
- **Enables traffic shaping** — Returns an immediate rejection when a virtual cluster's queue is full, so that you can redirect submissions to other virtual clusters instead of overwhelming a single virtual cluster.
- **Provides visibility** — Emits the per-virtual-cluster `JobsRunning` and `JobsInQueue` CloudWatch metrics in the `AWS/EMRContainers` namespace for active and in-queue job run counts every 5 minutes, which gives you a health signal for scheduling.

Getting started with concurrent job limits

You configure concurrent job limits with the `schedulerConfiguration` field on a virtual cluster. This field accepts two parameters:

`maxConcurrentJobRuns`

The maximum number of job runs that can be in the `RUNNING` state at any time.

`maxInQueueJobRuns`

The maximum number of job runs that can be in the `PENDING` or `SUBMITTED` state (queue depth) at any time.

AWS CLI

To set limits when you create a virtual cluster, specify `schedulerConfiguration` in your request.

```
aws emr-containers create-virtual-cluster \  
  --name my-virtual-cluster \  
  --container-provider '{ ... }' \  
  --scheduler-configuration '{  
    "maxConcurrentJobRuns": 500,  
    "maxInQueueJobRuns": 100  
  }'
```

To change limits on an existing virtual cluster, use the `update-virtual-cluster` command.

```
aws emr-containers update-virtual-cluster \  
  --id virtual-cluster-id \  
  --scheduler-configuration '{  
    "maxConcurrentJobRuns": 500,  
    "maxInQueueJobRuns": 100  
  }'
```

To remove the limits from a virtual cluster, pass an empty `schedulerConfiguration`. This clears the configuration, so no limits apply and the virtual cluster returns to default (unlimited) behavior. Note that *omitting* `schedulerConfiguration` from the request instead leaves the existing limits unchanged — you must pass an empty object to clear them.

```
aws emr-containers update-virtual-cluster \  
  --scheduler-configuration {}
```

```
--id virtual-cluster-id \  
--scheduler-configuration '{}'
```

To view the current limits and live job counts, use the `describe-virtual-cluster` command. The response includes both your `schedulerConfiguration` and a `SchedulerStatus` object with the current `activeJobRunCount` and `inQueueJobRunCount`.

Note

When you submit a job run to a virtual cluster whose queue is full, `StartJobRun` returns a `ValidationException`.

Choosing values for `maxConcurrentJobRuns` and `maxInQueueJobRuns`

The right limits depend on three things: how much work your Amazon EKS cluster can run at once, how bursty your submissions are, and how you want the virtual cluster to behave when it is full. Use the following guidance to pick a starting point, and then refine it from the live counters.

Setting `maxConcurrentJobRuns` (running slots)

`maxConcurrentJobRuns` is a job-granularity, count-based guardrail. A rough estimate here can protect the underlying Amazon EKS cluster from being degraded due to load and can improve availability.

- **Start from capacity divided by per-job footprint.** Base it on what each job *requests* (driver, executors, and memory overhead), and target approximately 70–80 percent of your namespace capacity to leave headroom for driver overhead, node scale-up, and bursts.
- **Cap the size of each job (T-shirt sizing).** Bound every job with `spark.dynamicAllocation.maxExecutors` and standardize on a few sizes — for example, Small (20 executors), Medium (100), and Large (approximately 500) — so that `maxConcurrentJobRuns` multiplied by the cap maps predictably to capacity instead of over-provisioning or under-provisioning for a variable average. For the cleanest math, route each size class to its own virtual cluster.
- **Tune from live counters.** Start conservative and raise the value gradually while you watch `activeJobRunCount` and the `JobsRunning` metric in the `AWS/EMRContainers` namespace.

Setting `maxInQueueJobRuns` (queue depth)

`maxInQueueJobRuns` controls how large a backlog the virtual cluster accepts before it starts rejecting submissions. It is a burst-absorption buffer. Consider the following factors.

- **Burst profile** — Size the queue to absorb the submission bursts that you expect above your running rate. If scheduled pipelines fire many jobs at once, a deeper queue prevents spurious rejections. Base the depth on your expected burst size rather than on a fixed multiple of `maxConcurrentJobRuns`, and validate it against the drain-time limit that follows.
- **Acceptable wait time** — Queued jobs wait for a running slot to free up. The job at the back of a full queue waits approximately the queue depth divided by the completion throughput. For example, if jobs finish at N per minute and the queue holds Q , the tail waits about Q divided by N minutes. Keep this within your SLA. Because buffered jobs fail after 30 minutes if no slot frees up, keep `maxInQueueJobRuns` small enough that a full queue drains well within 30 minutes at your steady completion rate. Otherwise, queued jobs time out.
- **Backpressure compared to buffering** — A deeper queue smooths bursts, but it delays the queue-full rejection that you use for traffic shaping and increases tail latency. A shallower queue fails fast, which gives clients an early, actionable signal to retry or route elsewhere. Choose based on whether you prefer to buffer load or shed and redirect it.
- **Client retry behavior** — When the queue is full, `StartJobRun` returns a `ValidationException`. Make sure that your submitters handle this exception — retry with backoff, or route the workload to another virtual cluster. Set the depth so that rejections occur only during genuine overload, not during routine operation.

Considerations for concurrent job limits

- No limits are applied by default. Existing virtual clusters and workloads are unaffected unless you explicitly set `schedulerConfiguration`.
- Because the counters are maintained across a distributed system, you can sometimes expect a small transient delta from the true value. Internal reconciliation corrects any drift.

Tutorials for Amazon EMR on EKS

This section describes common use cases for when you work with Amazon EMR on EKS applications. Each application is specialized and can take unique steps to configure. These topics provide instructions for using each application.

Topics

- [Using Delta Lake with Amazon EMR on EKS](#)
- [Using Apache Iceberg with Amazon EMR on EKS](#)
- [Using PyFlink](#)
- [Using AWS Glue with Flink](#)
- [Using Apache Hudi with Apache Flink](#)
- [Using cuDF Accelerator for Apache Spark with Amazon EMR on EKS](#)
- [Using Amazon Redshift integration for Apache Spark on Amazon EMR on EKS](#)
- [Using Volcano as a custom scheduler for Apache Spark on Amazon EMR on EKS](#)
- [Using YuniKorn as a custom scheduler for Apache Spark on Amazon EMR on EKS](#)

Using Delta Lake with Amazon EMR on EKS

Delta Lake is an open-source storage framework for building a Lakehouse architecture. The following shows how to set it up for use.

To use [Delta Lake](#) with Amazon EMR on EKS applications

1. When you start a job run to submit a Spark job in the application configuration, include the Delta Lake JAR files:

```
--job-driver '{"sparkSubmitJobDriver" : {  
    "sparkSubmitParameters" : "--jars local:///usr/share/aws/delta/lib/delta-  
core.jar,local:///usr/share/aws/delta/lib/delta-storage.jar,local:///usr/share/aws/  
delta/lib/delta-storage-s3-dynamodb.jar"}}'
```

Note

Amazon EMR releases 7.0.0 and higher uses Delta Lake 3.0, which renames `delta-core.jar` to `delta-spark.jar`. If you use Amazon EMR releases 7.0.0 or higher, be sure to use the correct file name, such as in the following example:

```
--jars local:///usr/share/aws/delta/lib/delta-spark.jar
```

2. Include Delta Lake additional configuration and use AWS Glue Data Catalog as your metastore.

```
--configuration-overrides '{
    "applicationConfiguration": [
    {
        "classification" : "spark-defaults",
        "properties" : {
            "spark.sql.extensions" : "io.delta.sql.DeltaSparkSessionExtension",

            "spark.sql.catalog.spark_catalog":"org.apache.spark.sql.delta.catalog.DeltaCatalog",
            "spark.hadoop.hive.metastore.client.factory.class":"com.amazonaws.glue.catalog.metastore.AWSGlueDataCatalogHiveClient"
        }
    }
    ]
}'
```

Using Apache Iceberg with Amazon EMR on EKS

The runtime JAR for Iceberg contains the necessary Iceberg classes for Spark runtime support. The following procedure shows how to start a job run using the Iceberg spark runtime.

To use Apache Iceberg with Amazon EMR on EKS applications

1. When you start a job run to submit a Spark job in the application configuration, include the Iceberg spark runtime JAR file:

```
--job-driver '{"sparkSubmitJobDriver" : {"sparkSubmitParameters" : "--jars
local:///usr/share/aws/iceberg/lib/iceberg-spark3-runtime.jar"}}'
```

2. Include Iceberg additional configuration:

```
--configuration-overrides '{
```

```

    "applicationConfiguration": [
      "classification" : "spark-defaults",
      "properties" : {
        "spark.sql.catalog.dev.warehouse" : "s3://amzn-s3-demo-bucket/EXAMPLE-
PREFIX/ ",
        "spark.sql.extensions ":"
org.apache.iceberg.spark.extensions.IcebergSparkSessionExtensions ",
        "spark.sql.catalog.dev" : "org.apache.iceberg.spark.SparkCatalog",
        "spark.sql.catalog.dev.catalog-impl" :
"org.apache.iceberg.aws.glue.GlueCatalog",
        "spark.sql.catalog.dev.io-impl": "org.apache.iceberg.aws.s3.S3FileIO"
      }
    ]
  }'

```

To learn more about Apache Iceberg release versions of EMR, see [Iceberg release history](#).

Spark session configurations for catalog integration

Spark session configurations for Iceberg AWS Glue catalog integration

This sample shows how to integrate Iceberg with the AWS Glue crawler:

```

spark-sql \
--conf spark.sql.catalog.rms = org.apache.iceberg.spark.SparkCatalog \
--conf spark.sql.catalog.rms.type = glue \
--conf spark.sql.catalog.rms.glue.id = glue RMS catalog ID \
--conf spark.sql.catalog.rms.glue.account-id = AWS account ID \

--conf spark.sql.extensions=
org.apache.iceberg.spark.extensions.IcebergSparkSessionExtensions

```

The following shows a sample query:

```
SELECT * FROM rms.rmsdb.table1
```

Spark session configurations for Iceberg REST AWS Glue catalog integration

This sample shows how to integrate Iceberg REST with the AWS Glue crawler:

```
spark-sql \
```

```
--conf spark.sql.catalog.rms = org.apache.iceberg.spark.SparkCatalog \  
--conf spark.sql.catalog.rms.type = rest \  
--conf spark.sql.catalog.rms.warehouse = glue RMS catalog ID \  
--conf spark.sql.catalog.rms.uri = glue endpoint URI/iceberg \  
--conf spark.sql.catalog.rms.rest.sigv4-enabled = true \  
--conf spark.sql.catalog.rms.rest.signing-name = glue \  
  
--conf spark.sql.extensions=  
    org.apache.iceberg.spark.extensions.IcebergSparkSessionExtensions
```

The following shows a sample query:

```
SELECT * FROM rms.rmsdb.table1
```

This configuration works for Redshift Managed Storage only. FGAC for Amazon S3 isn't supported.

Using PyFlink

Amazon EMR on EKS releases 6.15.0 and higher supports PyFlink. If you already have a PyFlink script, you can do one of the following:

- Create a custom image with your PyFlink script included.
- Upload your script to an Amazon S3 location

If you don't already have a script, you can use the following example to launch a PyFlink job. This example retrieves the script from S3. If you're using a custom image with your script already included in the image, you must update the script path to the location of where you stored your script. If the script is in an S3 location, Amazon EMR on EKS will retrieve the script and place it under the `/opt/flink/usrlib/` directory in the Flink container.

```
apiVersion: flink.apache.org/v1beta1  
kind: FlinkDeployment  
metadata:  
  name: python-example  
spec:  
  flinkVersion: v1_17  
  flinkConfiguration:  
    taskmanager.numberOfTaskSlots: "1"  
    executionRoleArn: job-execution-role  
    emrReleaseLabel: "emr-6.15.0-flink-latest"
```

```

jobManager:
  highAvailabilityEnabled: false
  replicas: 1
  resource:
    memory: "2048m"
    cpu: 1
taskManager:
  resource:
    memory: "2048m"
    cpu: 1
job:
  jarURI: s3://S3 bucket with your script/pyflink-script.py
  entryClass: "org.apache.flink.client.python.PythonDriver"
  args: ["-py", "/opt/flink/usrlib/pyflink-script.py"]
  parallelism: 1
  upgradeMode: stateless

```

Using AWS Glue with Flink

Amazon EMR on EKS with Apache Flink releases 6.15.0 and higher supports using the AWS Glue Data Catalog as a metadata store for streaming and batch SQL workflows.

You must first create an AWS Glue database named `default` that serves as your Flink SQL Catalog. This Flink Catalog stores metadata such as databases, tables, partitions, views, functions, and other information needed to access data in other external systems.

```

aws glue create-database \
  --database-input '{"Name":"default"}'

```

To enable AWS Glue support, use a `FlinkDeployment` spec. This example spec uses a Python script to quickly issue some Flink SQL statements to interact with the AWS Glue catalog.

```

apiVersion: flink.apache.org/v1beta1
kind: FlinkDeployment
metadata:
  name: python-example
spec:
  flinkVersion: v1_17
  flinkConfiguration:
    taskmanager.numberOfTaskSlots: "1"
    aws.glue.enabled: "true"

```

```

executionRoleArn: job-execution-role-arn;
emrReleaseLabel: "emr-6.15.0-flink-latest"
jobManager:
  highAvailabilityEnabled: false
  replicas: 1
  resource:
    memory: "2048m"
    cpu: 1
taskManager:
  resource:
    memory: "2048m"
    cpu: 1
job:
  jarURI: s3://<S3_bucket_with_your_script/pyflink-glue-script.py
  entryClass: "org.apache.flink.client.python.PythonDriver"
  args: ["-py", "/opt/flink/usrlib/pyflink-glue-script.py"]
  parallelism: 1
  upgradeMode: stateless

```

The following is an example of what your PyFlink script might look like.

```

import logging
import sys
from pyflink.datastream import StreamExecutionEnvironment
from pyflink.table import StreamTableEnvironment

def glue_demo():
    env = StreamExecutionEnvironment.get_execution_environment()
    t_env = StreamTableEnvironment.create(stream_execution_environment=env)
    t_env.execute_sql("""
        CREATE CATALOG glue_catalog WITH (
            'type' = 'hive',
            'default-database' = 'default',
            'hive-conf-dir' = '/glue/confs/hive/conf',
            'hadoop-conf-dir' = '/glue/confs/hadoop/conf'
        )
    """)
    t_env.execute_sql("""
        USE CATALOG glue_catalog;
    """)
    t_env.execute_sql("""
        DROP DATABASE IF EXISTS eks_flink_db CASCADE;
    """)

```

```

t_env.execute_sql("""
    CREATE DATABASE IF NOT EXISTS eks_flink_db WITH ('hive.database.location-
uri'= 's3a://S3-bucket-to-store-metadata/flink/flink-glue-for-hive/warehouse/');
    """)
t_env.execute_sql("""
    USE eks_flink_db;
    """)
t_env.execute_sql("""
    CREATE TABLE IF NOT EXISTS eksglueorders (
        order_number BIGINT,
        price          DECIMAL(32,2),
        buyer          ROW first_name STRING, last_name STRING,
        order_time      TIMESTAMP(3)
    ) WITH (
        'connector' = 'datagen'
    );
    """)
t_env.execute_sql("""
    CREATE TABLE IF NOT EXISTS eksdestglueorders (
        order_number BIGINT,
        price          DECIMAL(32,2),
        buyer          ROW first_name STRING, last_name STRING,
        order_time      TIMESTAMP(3)
    ) WITH (
        'connector' = 'filesystem',
        'path' = 's3://S3-bucket-to-store-metadata/flink/flink-glue-for-hive/
warehouse/eksdestglueorders',
        'format' = 'json'
    );
    """)
t_env.execute_sql("""
    CREATE TABLE IF NOT EXISTS print_table (
        order_number BIGINT,
        price          DECIMAL(32,2),
        buyer          ROW first_name STRING, last_name STRING,
        order_time      TIMESTAMP(3)
    ) WITH (
        'connector' = 'print'
    );
    """)
t_env.execute_sql("""
    EXECUTE STATEMENT SET
    BEGIN
    INSERT INTO eksdestglueorders SELECT * FROM eksglueorders LIMIT 10;

```

```
INSERT INTO print_table SELECT * FROM eksdestglueorders;
END;
""")

if __name__ == '__main__':
    logging.basicConfig(stream=sys.stdout, level=logging.INFO, format="%(message)s")
    glue_demo()
```

Using Apache Hudi with Apache Flink

Apache Hudi is an open-source data management framework with record-level operations such as insert, update, upsert, and delete that you can use to simplify data management and data pipeline development. When combined with efficient data management in Amazon S3, Hudi lets you ingest and update data in real time. Hudi maintains metadata of all of the operations that you run on the dataset, so all of the actions remain atomic and consistent.

Apache Hudi is available on Amazon EMR on EKS with Apache Flink with Amazon EMR releases 7.2.0 and higher. See the following steps to learn how to get started and submit Apache Hudi jobs.

Submit an Apache Hudi job

See the following steps to learn how to submit an Apache Hudi job.

1. Create an AWS Glue database named default.

```
aws glue create-database --database-input '{"Name":"default"}'
```

2. Follow the [Flink Kubernetes Operator SQL Example](#) to build the `flink-sql-runner.jar` file.
3. Create a Hudi SQL script like the following.

```
CREATE CATALOG hudi_glue_catalog WITH (
  'type' = 'hudi',
  'mode' = 'hms',
  'table.external' = 'true',
  'default-database' = 'default',
  'hive.conf.dir' = '/glue/confs/hive/conf/',
  'catalog.path' = 's3://<hudi-example-bucket>/FLINK_HUDI/warehouse/'
);
```



```

USE CATALOG hudi_glue_catalog;
CREATE DATABASE IF NOT EXISTS hudi_db;
use hudi_db;

CREATE TABLE IF NOT EXISTS hudi-flink-example-table(
    uuid VARCHAR(20),
    name VARCHAR(10),
    age INT,
    ts TIMESTAMP(3),
    `partition` VARCHAR(20)
)
PARTITIONED BY (`partition`)
WITH (
    'connector' = 'hudi',
    'path' = 's3://<hudi-example-bucket>/hudi-flink-example-table',
    'hive_sync.enable' = 'true',
    'hive_sync.mode' = 'glue',
    'hive_sync.table' = 'hudi-flink-example-table',
    'hive_sync.db' = 'hudi_db',
    'compaction.delta_commits' = '1',
    'hive_sync.partition_fields' = 'partition',
    'hive_sync.partition_extractor_class' =
    'org.apache.hudi.hive.MultiPartKeysValueExtractor',
    'table.type' = 'COPY_ON_WRITE'
);

EXECUTE STATEMENT SET
BEGIN

INSERT INTO hudi-flink-example-table VALUES
    ('id1','Alex',23,TIMESTAMP '1970-01-01 00:00:01','par1'),
    ('id2','Stephen',33,TIMESTAMP '1970-01-01 00:00:02','par1'),
    ('id3','Julian',53,TIMESTAMP '1970-01-01 00:00:03','par2'),
    ('id4','Fabian',31,TIMESTAMP '1970-01-01 00:00:04','par2'),
    ('id5','Sophia',18,TIMESTAMP '1970-01-01 00:00:05','par3'),
    ('id6','Emma',20,TIMESTAMP '1970-01-01 00:00:06','par3'),
    ('id7','Bob',44,TIMESTAMP '1970-01-01 00:00:07','par4'),
    ('id8','Han',56,TIMESTAMP '1970-01-01 00:00:08','par4');

END;

```

4. Upload your Hudi SQL script and the `flink-sql-runner.jar` file to an S3 location.
5. In your `FlinkDeployments` YAML file, set `hudi.enabled` to `true`.

```
spec:
  flinkConfiguration:
    hudi.enabled: "true"
```

6. Create a YAML file to run your configuration. This example file is named `hudi-write.yaml`.

```
apiVersion: flink.apache.org/v1beta1
kind: FlinkDeployment
metadata:
  name: hudi-write-example
spec:
  flinkVersion: v1_18
  flinkConfiguration:
    taskmanager.numberOfTaskSlots: "2"
    hudi.enabled: "true"
  executionRoleArn: "<JobExecutionRole>"
  emrReleaseLabel: "emr-7.13.0-flink-latest"
  jobManager:
    highAvailabilityEnabled: false
    replicas: 1
    resource:
      memory: "2048m"
      cpu: 1
  taskManager:
    resource:
      memory: "2048m"
      cpu: 1
  job:
    jarURI: local:///opt/flink/usrlib/flink-sql-runner.jar
    args: ["/opt/flink/scripts/hudi-write.sql"]
    parallelism: 1
    upgradeMode: stateless
  podTemplate:
    spec:
      initContainers:
        - name: flink-sql-script-download
          args:
            - s3
            - cp
            - s3://<s3_location>/hudi-write.sql
            - /flink-scripts
      image: amazon/aws-cli:latest
```

```

    imagePullPolicy: Always
    resources: {}
    terminationMessagePath: /dev/termination-log
    terminationMessagePolicy: File
    volumeMounts:
      - mountPath: /flink-scripts
        name: flink-scripts
- name: flink-sql-runner-download
  args:
    - s3
    - cp
    - s3://<s3_location>/flink-sql-runner.jar
    - /flink-artifacts
  image: amazon/aws-cli:latest
  imagePullPolicy: Always
  resources: {}
  terminationMessagePath: /dev/termination-log
  terminationMessagePolicy: File
  volumeMounts:
    - mountPath: /flink-artifacts
      name: flink-artifact
containers:
- name: flink-main-container
  volumeMounts:
    - mountPath: /opt/flink/scripts
      name: flink-scripts
    - mountPath: /opt/flink/usrlib
      name: flink-artifact
volumes:
- emptyDir: {}
  name: flink-scripts
- emptyDir: {}
  name: flink-artifact

```

7. Submit a Flink Hudi job to the [Flink Kubernetes operator](#).

```
kubectl apply -f hudi-write.yaml
```

Using cuDF Accelerator for Apache Spark with Amazon EMR on EKS

With Amazon EMR on EKS, you can run jobs for the NVIDIA cuDF Accelerator for Apache Spark. This tutorial covers how to run Spark jobs using cuDF on EC2 graphics processing unit (GPU) instance types. The tutorial uses the following versions:

- Amazon EMR on EKS release version 6.9.0 and later
- Apache Spark 3.x

You can accelerate Spark with Amazon EC2 GPU instance types by using the NVIDIA [cuDF Accelerator for Apache Spark](#) plugin. When you use these technologies together, you accelerate your data science pipelines without having to make any code changes. This reduces the run time needed for data processing and model training. By getting more done in less time, you spend less on the cost of infrastructure.

Before you begin, make sure you have the following resources.

- Amazon EMR on EKS virtual cluster
- Amazon EKS cluster with a GPU enabled node group

An Amazon EKS virtual cluster is a registered handle to the Kubernetes namespace on an Amazon EKS cluster, and is managed by Amazon EMR on EKS. The handle allows Amazon EMR to use the Kubernetes namespace as a destination for running jobs. For more information on how to set up a virtual cluster, see [Setting up Amazon EMR on EKS](#) in this guide.

You must configure the Amazon EKS virtual cluster with a node group that has GPU instances. You must configure the nodes with an Nvidia device plugin. See [managed node groups](#) to learn more.

To configure your Amazon EKS cluster to add GPU-enabled node groups, perform the following procedure:

To add GPU enabled node groups

1. Create a GPU-enabled node group with the following [create-nodegroup](#) command. Be sure to substitute the correct parameters for your Amazon EKS cluster. Use an instance type that supports Spark cuDF, such as P4, P3, G5 or G4dn.

```
aws eks create-nodegroup \
  --cluster-name EKS_CLUSTER_NAME \
  --nodegroup-name NODEGROUP_NAME \
  --scaling-config minSize=0,maxSize=5,desiredSize=2 CHOOSE_APPROPRIATELY \
  --ami-type AL2_x86_64_GPU \
  --node-role NODE_ROLE \
  --subnets SUBNETS_SPACE_DELIMITED \
  --remote-access ec2SshKey= SSH_KEY \
  --instance-types GPU_INSTANCE_TYPE \
  --disk-size DISK_SIZE \
  --region AWS_REGION
```

2. Install the Nvidia device plugin in your cluster to emit the number of GPUs on each node of your cluster and to run GPU-enabled containers in your cluster. Run the following code to install the plugin:

```
kubectl apply -f https://raw.githubusercontent.com/NVIDIA/k8s-device-plugin/v0.9.0/nvidia-device-plugin.yml
```

3. To validate how many GPUs are available on each node of your cluster, run the following command:

```
kubectl get nodes "-o=custom-  
columns=NAME:.metadata.name,GPU:.status.allocatable.nvidia\.com/gpu"
```

To run a Spark cuDF job

1. Submit a Spark cuDF job to your Amazon EMR on EKS cluster. The following code shows an example of a command to start the job. The first time you run the job, it might take a few minutes to download the image and cache it on the node.

```
aws emr-containers start-job-run \
  --virtual-cluster-id VIRTUAL_CLUSTER_ID \
  --execution-role-arn JOB_EXECUTION_ROLE \
  --release-label emr-6.9.0-spark-rapids-latest \
  --job-driver '{"sparkSubmitJobDriver": {"entryPoint": "local:///usr/lib/  
spark/examples/jars/spark-examples.jar","entryPointArguments": ["10000"],  
  "sparkSubmitParameters":"--class org.apache.spark.examples.SparkPi "}}' \
  ---configuration-overrides '{"applicationConfiguration": [{"classification":  
  "spark-defaults","properties": {"spark.executor.instances":
```

```
"2","spark.executor.memory": "2G"}]], "monitoringConfiguration":
{"cloudWatchMonitoringConfiguration": {"logGroupName": "LOG_GROUP
_NAME"}, "s3MonitoringConfiguration": {"logUri": "LOG_GROUP_STREAM"}}}'
```

2. To validate that the Spark cuDF Accelerator is enabled, check the Spark driver logs. These logs are stored either in CloudWatch or in the S3 location you specify when you run the `start-job-run` command. The following example generally shows what the log lines look like:

```
22/11/15 00:12:44 INFO RapidsPluginUtils: RAPIDS Accelerator build:
{version=22.08.0-amzn-0, user=release, url=, date=2022-11-03T03:32:45Z, revision=,
cudf_version=22.08.0, branch=}
22/11/15 00:12:44 INFO RapidsPluginUtils: RAPIDS Accelerator JNI build:
{version=22.08.0, user=, url=https://github.com/NVIDIA/spark-rapids-jni.git,
date=2022-08-18T04:14:34Z, revision=a1b23cd_sample, branch=HEAD}
22/11/15 00:12:44 INFO RapidsPluginUtils: cudf build: {version=22.08.0,
user=, url=https://github.com/rapidsai/cudf.git, date=2022-08-18T04:14:34Z,
revision=a1b23ce_sample, branch=HEAD}
22/11/15 00:12:44 WARN RapidsPluginUtils: RAPIDS Accelerator 22.08.0-amzn-0 using
cudf 22.08.0.
22/11/15 00:12:44 WARN RapidsPluginUtils:
spark.rapids.sql.multiThreadedRead.numThreads is set to 20.
22/11/15 00:12:44 WARN RapidsPluginUtils: RAPIDS Accelerator is enabled, to disable
GPU support set `spark.rapids.sql.enabled` to false.
22/11/15 00:12:44 WARN RapidsPluginUtils: spark.rapids.sql.explain is set to
`NOT_ON_GPU`. Set it to 'NONE' to suppress the diagnostics logging about the query
placement on the GPU.
```

3. To see the operations that will be run on a GPU, perform the following steps to enable extra logging. Note the `"spark.rapids.sql.explain : ALL"` config.

```
aws emr-containers start-job-run \
--virtual-cluster-id VIRTUAL_CLUSTER_ID \
--execution-role-arn JOB_EXECUTION_ROLE \
--release-label emr-6.9.0-spark-rapids-latest \
--job-driver '{"sparkSubmitJobDriver": {"entryPoint": "local:///usr/lib/
spark/examples/jars/spark-examples.jar", "entryPointArguments": ["10000"],
"sparkSubmitParameters": "--class org.apache.spark.examples.SparkPi "}}' \
---configuration-overrides '{"applicationConfiguration":
[{"classification": "spark-defaults", "properties":
{"spark.rapids.sql.explain": "ALL", "spark.executor.instances":
"2", "spark.executor.memory": "2G"}]], "monitoringConfiguration":
```

```
{ "cloudWatchMonitoringConfiguration": { "logGroupName":
  "LOG_GROUP_NAME"}, "s3MonitoringConfiguration": { "logUri": "LOG_GROUP_STREAM"}}}'
```

The previous command is an example of a job that uses the GPU. Its output would look something like the example below. Refer to this key for help to understand the output:

- * – marks an operation that works on a GPU
- ! – marks an operation that can't run on a GPU
- @ – marks an operation that works on a GPU, but won't get to run because it's inside a plan that can't run on a GPU

```
22/11/15 01:22:58 INFO GpuOverrides: Plan conversion to the GPU took 118.64 ms
22/11/15 01:22:58 INFO GpuOverrides: Plan conversion to the GPU took 4.20 ms
22/11/15 01:22:58 INFO GpuOverrides: GPU plan transition optimization took 8.37 ms
22/11/15 01:22:59 WARN GpuOverrides:
  *Exec <ProjectExec> will run on GPU
    *Expression <Alias> substring(cast(date#149 as string), 0, 7) AS month#310
will run on GPU
    *Expression <Substring> substring(cast(date#149 as string), 0, 7) will run
on GPU
      *Expression <Cast> cast(date#149 as string) will run on GPU
    *Exec <SortExec> will run on GPU
      *Expression <SortOrder> date#149 ASC NULLS FIRST will run on GPU
    *Exec <ShuffleExchangeExec> will run on GPU
      *Partitioning <RangePartitioning> will run on GPU
        *Expression <SortOrder> date#149 ASC NULLS FIRST will run on GPU
    *Exec <UnionExec> will run on GPU
      !Exec <ProjectExec> cannot run on GPU because not all expressions can
be replaced
        @Expression <AttributeReference> customerID#0 could run on GPU
        @Expression <Alias> Charge AS kind#126 could run on GPU
          @Expression <Literal> Charge could run on GPU
          @Expression <AttributeReference> value#129 could run on GPU
          @Expression <Alias> add_months(2022-11-15, cast(-(cast(_we0#142 as
bigint) + last_month#128L) as int)) AS date#149 could run on GPU
            ! <AddMonths> add_months(2022-11-15, cast(-
(cast(_we0#142 as bigint) + last_month#128L) as int)) cannot run
on GPU because GPU does not currently support the operator class
org.apache.spark.sql.catalyst.expressions.AddMonths
              @Expression <Literal> 2022-11-15 could run on GPU
```

```

        @Expression <Cast> cast(-(cast(_we0#142 as bigint) +
last_month#128L) as int) could run on GPU
        @Expression <UnaryMinus> -(cast(_we0#142 as bigint) +
last_month#128L) could run on GPU
        @Expression <Add> (cast(_we0#142 as bigint) +
last_month#128L) could run on GPU
        @Expression <Cast> cast(_we0#142 as bigint) could run on
GPU
        @Expression <AttributeReference> _we0#142 could run on
GPU
        @Expression <AttributeReference> last_month#128L could run
on GPU

```

Using Amazon Redshift integration for Apache Spark on Amazon EMR on EKS

With Amazon EMR release 6.9.0 and later, every release image includes a connector between [Apache Spark](#) and Amazon Redshift. This way, you can use Spark on Amazon EMR on EKS to process data stored in Amazon Redshift. The integration is based on the [spark-redshift open-source connector](#). For Amazon EMR on EKS, the [Amazon Redshift integration for Apache Spark](#) is included as a native integration.

Topics

- [Launching a Spark application using the Amazon Redshift integration for Apache Spark](#)
- [Authenticating with the Amazon Redshift integration for Apache Spark](#)
- [Reading and writing from and to Amazon Redshift](#)
- [Considerations and limitations when using the Spark connector](#)

Launching a Spark application using the Amazon Redshift integration for Apache Spark

To use the integration, you must pass the required Spark Redshift dependencies with your Spark job. You must use `--jars` to include Redshift connector-related libraries. To see other file locations supported by the `--jars` option, see the [Advanced Dependency Management](#) section of the Apache Spark documentation.

- spark-redshift.jar
- spark-avro.jar
- RedshiftJDBC.jar
- minimal-json.jar

To launch a Spark application with the Amazon Redshift integration for Apache Spark on Amazon EMR on EKS release 6.9.0 or later, use the following example command. Note that the paths listed with the `--conf spark.jars` option are the default paths for the JAR files.

```
aws emr-containers start-job-run \  
  
--virtual-cluster-id cluster_id \  
--execution-role-arn arn \  
--release-label emr-6.9.0-latest \  
--job-driver '{  
    "sparkSubmitJobDriver": {  
        "entryPoint": "s3://script_path",  
        "sparkSubmitParameters":  
            "--conf spark.kubernetes.file.upload.path=s3://upload_path  
            --conf spark.jars=  
                /usr/share/aws/redshift/jdbc/RedshiftJDBC.jar,  
                /usr/share/aws/redshift/spark-redshift/lib/spark-redshift.jar,  
                /usr/share/aws/redshift/spark-redshift/lib/spark-avro.jar,  
                /usr/share/aws/redshift/spark-redshift/lib/minimal-json.jar"  
    }  
}'
```

Authenticating with the Amazon Redshift integration for Apache Spark

The following sections show authentication options with Amazon Redshift when you're integrating with Apache Spark. The sections show how to retrieve login credentials and also details regarding using the JDBC driver with IAM authentication.

Use AWS Secrets Manager to retrieve credentials and connect to Amazon Redshift

You can store credentials in Secrets Manager to authenticate securely to Amazon Redshift. You can have your Spark job call the `GetSecretValue` API to fetch the credentials:

```
from pyspark.sql import SQLContextimport boto3

sc = # existing SparkContext
sql_context = SQLContext(sc)

secretsmanager_client = boto3.client('secretsmanager',
    region_name=os.getenv('AWS_REGION'))
secret_manager_response = secretsmanager_client.get_secret_value(
    SecretId='string',
    VersionId='string',
    VersionStage='string'
)
username = # get username from secret_manager_response
password = # get password from secret_manager_response
url = "jdbc:redshift://redshifthost:5439/database?user=" + username + "&password="
    + password

# Access to Redshift cluster using Spark
```

Use IAM based authentication with Amazon EMR on EKS job execution role

Starting with Amazon EMR on EKS release 6.9.0, the Amazon Redshift JDBC driver version 2.1 or higher is packaged into the environment. With JDBC driver 2.1 and higher, you can specify the JDBC URL and not include the raw username and password. Instead, you can specify `jdbc:redshift:iam://` scheme. This commands the JDBC driver to use your Amazon EMR on EKS job execution role to fetch the credentials automatically.

See [Configure a JDBC or ODBC connection to use IAM credentials](#) in the *Amazon Redshift Management Guide* for more information.

The following example URL uses a `jdbc:redshift:iam://` scheme.

```
jdbc:redshift:iam://examplecluster.abc123xyz789.us-west-2.redshift.amazonaws.com:5439/
dev
```

The following permissions are required for your job execution role when it meets the provided conditions.

Permission	Conditions when required for job execution role
<code>redshift:GetClusterCredentials</code>	Required for JDBC driver to fetch the credentials from Amazon Redshift
<code>redshift:DescribeCluster</code>	Required if you specify the Amazon Redshift cluster and AWS Region in the JDBC URL instead of endpoint
<code>redshift-serverless:GetCredentials</code>	Required for JDBC driver to fetch the credentials from Amazon Redshift Serverless
<code>redshift-serverless:GetWorkgroup</code>	Required if you are using Amazon Redshift Serverless and you specify the URL in terms of workgroup name and Region

Your job execution role policy should have the following permissions.

```
{
    "Effect": "Allow",
    "Action": [
        "redshift:GetClusterCredentials",
        "redshift:DescribeCluster",
        "redshift-serverless:GetCredentials",
        "redshift-serverless:GetWorkgroup"
    ],
    "Resource": [
        "arn:aws:redshift:AWS_REGION:ACCOUNT_ID:dbname:CLUSTER_NAME/DATABASE_NAME",
        "arn:aws:redshift:AWS_REGION:ACCOUNT_ID:dbuser:DATABASE_NAME/USER_NAME"
    ]
}
```

Authenticate to Amazon Redshift with a JDBC driver

Set username and password inside the JDBC URL

To authenticate a Spark job to an Amazon Redshift cluster, you can specify the Amazon Redshift database name and password in the JDBC URL.

Note

If you pass the database credentials in the URL, anyone who has access to the URL can also access the credentials. This method isn't generally recommended because it's not a secure option.

If security isn't a concern for your application, you can use the following format to set the username and password in the JDBC URL:

```
jdbc:redshift://redshifthost:5439/database?user=username&password=password
```

Reading and writing from and to Amazon Redshift

The following code examples use PySpark to read and write sample data from and to an Amazon Redshift database with a data source API and with SparkSQL.

Data source API

Use PySpark to read and write sample data from and to an Amazon Redshift database with a data source API.

```
import boto3
from pyspark.sql import SQLContext

sc = # existing SparkContext
sql_context = SQLContext(sc)

url = "jdbc:redshift:iam://redshifthost:5439/database"
aws_iam_role_arn = "arn:aws:iam::accountID:role/roleName"

df = sql_context.read \
    .format("io.github.spark_redshift_community.spark.redshift") \
    .option("url", url) \
    .option("dbtable", "tableName") \
    .option("tempdir", "s3://path/for/temp/data") \
    .option("aws_iam_role", "aws_iam_role_arn") \
    .load()

df.write \
    .format("io.github.spark_redshift_community.spark.redshift") \
```

```
.option("url", url) \
.option("dbtable", "tableName_copy") \
.option("tempdir", "s3://path/for/temp/data") \
.option("aws_iam_role", "aws_iam_role_arn") \
.mode("error") \
.save()
```

SparkSQL

Use PySpark to read and write sample data from and to an Amazon Redshift database using SparkSQL.

```
import boto3
import json
import sys
import os
from pyspark.sql import SparkSession

spark = SparkSession \
    .builder \
    .enableHiveSupport() \
    .getOrCreate()

url = "jdbc:redshift:iam://redshifthost:5439/database"
aws_iam_role_arn = "arn:aws:iam::accountID:role/roleName"

bucket = "s3://path/for/temp/data"
tableName = "tableName" # Redshift table name

s = f"""CREATE TABLE IF NOT EXISTS {tableName} (country string, data string)
    USING io.github.spark_redshift_community.spark.redshift
    OPTIONS (dbtable '{tableName}', tempdir '{bucket}', url '{url}', aws_iam_role
    '{aws_iam_role_arn'} '); """

spark.sql(s)

columns = ["country" ,"data"]
data = [("test-country", "test-data")]
df = spark.sparkContext.parallelize(data).toDF(columns)

# Insert data into table
df.write.insertInto(tableName, overwrite=False)
df = spark.sql(f"SELECT * FROM {tableName}")
```

```
df.show()
```

Considerations and limitations when using the Spark connector

The Spark connector supports a variety of ways to manage credentials, to configure security, and to connect with other AWS services. Get familiar with the recommendations in this list in order to configure a functional and resilient connection.

- We recommend that you activate SSL for the JDBC connection from Spark on Amazon EMR to Amazon Redshift.
- We recommend that you manage the credentials for the Amazon Redshift cluster in AWS Secrets Manager as a best practice. See [Using AWS Secrets Manager to retrieve credentials for connecting to Amazon Redshift](#) for an example.
- We recommend that you pass an IAM role with the parameter `aws_iam_role` for the Amazon Redshift authentication parameter.
- The parameter `tempformat` currently doesn't support the Parquet format.
- The `tempdir` URI points to an Amazon S3 location. This temp directory isn't cleaned up automatically and therefore could add additional cost.
- Consider the following recommendations for Amazon Redshift:
 - We recommend that you block public access to the Amazon Redshift cluster.
 - We recommend that you turn on [Amazon Redshift audit logging](#).
 - We recommend turn on [Amazon Redshift at-rest encryption](#).
- Consider the following recommendations for Amazon S3:
 - We recommend [blocking public access to Amazon S3 buckets](#).
 - We recommend that you use [Amazon S3 server-side encryption](#) to encrypt the S3 buckets that you use.
 - We recommend that you use [Amazon S3 lifecycle policies](#) to define the retention rules for the S3 bucket.
 - Amazon EMR always verifies code imported from open-source into the image. For security, we don't support encoding AWS access keys in the `tempdir` URI as an authentication method from Spark to Amazon S3.

For more information on using the connector and its supported parameters, see the following resources:

- [Amazon Redshift integration for Apache Spark](#) in the *Amazon Redshift Management Guide*
- The [spark-redshift community repository](#) on Github

Using Volcano as a custom scheduler for Apache Spark on Amazon EMR on EKS

With Amazon EMR on EKS, you can use Spark operator or spark-submit to run Spark jobs with Kubernetes custom schedulers. This tutorial covers how to run Spark jobs with a Volcano scheduler on a custom queue.

Overview

[Volcano](#) can help manage Spark scheduling with advanced functions such as queue scheduling, fair-share scheduling, and resource reservation. For more information on the benefits of Volcano, see [Why Spark chooses Volcano as built-in batch scheduler on Kubernetes](#) on The Linux Foundation's *CNCF blog*.

Install and set up Volcano

1. Choose one of the following kubectl commands to install Volcano, depending on your architectural needs:

```
# x86_64
kubectl apply -f https://raw.githubusercontent.com/volcano-sh/volcano/v1.5.1/
installer/volcano-development.yaml
# arm64:
kubectl apply -f https://raw.githubusercontent.com/volcano-sh/volcano/v1.5.1/
installer/volcano-development-arm64.yaml
```

2. Prepare a sample Volcano queue. A queue is a collection of [PodGroups](#). The queue adopts FIFO and is the basis for resource division.

```
cat << EOF > volcanoQ.yaml
apiVersion: scheduling.volcano.sh/v1beta1
kind: Queue
```

```

metadata:
  name: sparkqueue
spec:
  weight: 4
  reclaimable: false
  capability:
    cpu: 10
    memory: 20Gi
EOF

kubectl apply -f volcanoQ.yaml

```

3. Upload a sample PodGroup manifest to Amazon S3. PodGroup is a group of pods with strong association. You typically use a PodGroup for batch scheduling. Submit the following sample PodGroup to the queue that you defined in the previous step.

```

cat << EOF > podGroup.yaml
apiVersion: scheduling.volcano.sh/v1beta1
kind: PodGroup
spec:
  # Set minMember to 1 to make a driver pod
  minMember: 1
  # Specify minResources to support resource reservation.
  # Consider the driver pod resource and executors pod resource.
  # The available resources should meet the minimum requirements of the Spark job
  # to avoid a situation where drivers are scheduled, but they can't schedule
  # sufficient executors to progress.
  minResources:
    cpu: "1"
    memory: "1Gi"
  # Specify the queue. This defines the resource queue that the job should be
  # submitted to.
  queue: sparkqueue
EOF

aws s3 mv podGroup.yaml s3://bucket-name

```

Run a Spark application with Volcano scheduler with the Spark operator

1. If you haven't already, complete the steps in the following sections to get set up:

- a. [Install and set up Volcano](#)
- b. [Setting up the Spark operator for Amazon EMR on EKS](#)
- c. [Install the Spark operator](#)

Include the following arguments when you run the `helm install spark-operator-demo` command:

```
--set batchScheduler.enable=true  
--set webhook.enable=true
```

2. Create a SparkApplication definition file `spark-pi.yaml` with `batchScheduler` configured.

```
apiVersion: "sparkoperator.k8s.io/v1beta2"  
kind: SparkApplication  
metadata:  
  name: spark-pi  
  namespace: spark-operator  
spec:  
  type: Scala  
  mode: cluster  
  image: "895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/emr-6.10.0:latest"  
  imagePullPolicy: Always  
  mainClass: org.apache.spark.examples.SparkPi  
  mainApplicationFile: "local:///usr/lib/spark/examples/jars/spark-examples.jar"  
  sparkVersion: "3.3.1"  
  batchScheduler: "volcano" #Note: You must specify the batch scheduler name as  
'volcano'  
  restartPolicy:  
    type: Never  
  volumes:  
    - name: "test-volume"  
      hostPath:  
        path: "/tmp"  
        type: Directory  
  driver:  
    cores: 1  
    coreLimit: "1200m"  
    memory: "512m"  
    labels:  
      version: 3.3.1
```

```

serviceAccount: emr-containers-sa-spark
volumeMounts:
  - name: "test-volume"
    mountPath: "/tmp"
executor:
  cores: 1
  instances: 1
  memory: "512m"
  labels:
    version: 3.3.1
  volumeMounts:
    - name: "test-volume"
      mountPath: "/tmp"

```

3. Submit the Spark application with the following command. This also creates a SparkApplication object called spark-pi:

```
kubectl apply -f spark-pi.yaml
```

4. Check events for the SparkApplication object with the following command:

```
kubectl describe pods spark-pi-driver --namespace spark-operator
```

The first pod event will show that Volcano has scheduled the pods:

Type	Reason	Age	From	Message
----	-----	----	----	-----
Normal	Scheduled	23s	volcano	Successfully assigned default/spark-pi-driver to integration-worker2

Run a Spark application with Volcano scheduler with spark-submit

1. First, complete the steps in the [Setting up spark-submit for Amazon EMR on EKS](#) section. You must build your spark-submit distribution with Volcano support. For more information, see the **Build section** of [Using Volcano as Customized Scheduler for Spark on Kubernetes](#) in the *Apache Spark documentation*.
2. Set the values for the following environment variables:

```
export SPARK_HOME=spark-home
```

```
export MASTER_URL=k8s://Amazon-EKS-cluster-endpoint
```

3. Submit the Spark application with the following command:

```
$SPARK_HOME/bin/spark-submit \
--class org.apache.spark.examples.SparkPi \
--master $MASTER_URL \
--conf spark.kubernetes.container.image=895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/emr-6.10.0:latest \
--conf spark.kubernetes.authenticate.driver.serviceAccountName=spark \
--deploy-mode cluster \
--conf spark.kubernetes.namespace=spark-operator \
--conf spark.kubernetes.scheduler.name=volcano \
--conf spark.kubernetes.scheduler.volcano.podGroupTemplateFile=/path/to/podgroup-
template.yaml \
--conf
spark.kubernetes.driver.pod.featureSteps=org.apache.spark.deploy.k8s.features.VolcanoFeatu
\
--conf
spark.kubernetes.executor.pod.featureSteps=org.apache.spark.deploy.k8s.features.VolcanoFea
\
local:///usr/lib/spark/examples/jars/spark-examples.jar 20
```

4. Check events for the SparkApplication object with the following command:

```
kubectl describe pod spark-pi --namespace spark-operator
```

The first pod event will show that Volcano has scheduled the pods:

Type	Reason	Age	From	Message
----	-----	----	----	-----
Normal	Scheduled	23s	volcano	Successfully assigned default/spark-pi-driver to integration-worker2

Using YuniKorn as a custom scheduler for Apache Spark on Amazon EMR on EKS

With Amazon EMR on EKS, you can use Spark operator or spark-submit to run Spark jobs with Kubernetes custom schedulers. This tutorial covers how to run Spark jobs with a YuniKorn scheduler on a custom queue and gang scheduling.

Overview

[Apache YuniKorn](#) can help manage Spark scheduling with app-aware scheduling so that you can have fine-grained control on resource quotas and priorities. With gang scheduling, YuniKorn schedules an app only when the minimal resource request for the app can be satisfied. For more information, see [What is gang scheduling](#) on the Apache YuniKorn documentation site.

Create your cluster and get set up for YuniKorn

Use the following steps to deploy an Amazon EKS cluster. You can change the AWS Region (region) and Availability Zones (availabilityZones).

1. Define the Amazon EKS cluster:

```
cat <<EOF >eks-cluster.yaml
---
apiVersion: eksctl.io/v1alpha5
kind: ClusterConfig

metadata:
  name: emr-eks-cluster
  region: eu-west-1

vpc:
  clusterEndpoints:
    publicAccess: true
    privateAccess: true

iam:
  withOIDC: true

nodeGroups:
  - name: spark-jobs
    labels: { app: spark }
    instanceType: m5.xlarge
    desiredCapacity: 2
    minSize: 2
    maxSize: 3
    availabilityZones: ["eu-west-1a"]
EOF
```

2. Create the cluster:

```
eksctl create cluster -f eks-cluster.yaml
```

3. Create the namespace `spark-job` where you will execute the Spark job:

```
kubectl create namespace spark-job
```

4. Next, create a Kubernetes role and role binding. This is required for the service account that the Spark job run uses.
 - a. Define the service account, role, and role binding for Spark jobs.

```
cat <<EOF >emr-job-execution-rbac.yaml
---
apiVersion: v1
kind: ServiceAccount
metadata:
  name: spark-sa
  namespace: spark-job
automountServiceAccountToken: false
---
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: spark-role
  namespace: spark-job
rules:
  - apiGroups: ["", "batch", "extensions"]
    resources: ["configmaps", "serviceaccounts", "events", "pods", "pods/
exec", "pods/log", "pods/
portforward", "secrets", "services", "persistentvolumeclaims"]
    verbs: ["create", "delete", "get", "list", "patch", "update", "watch"]
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: spark-sa-rb
  namespace: spark-job
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: Role
  name: spark-role
subjects:
```

```
- kind: ServiceAccount
  name: spark-sa
  namespace: spark-job
EOF
```

- b. Apply the Kubernetes role and role binding definition with the following command:

```
kubectl apply -f emr-job-execution-rbac.yaml
```

Install and set up YuniKorn

1. Use the following `kubectl` command to create a namespace `yunikorn` to deploy the YuniKorn scheduler:

```
kubectl create namespace yunikorn
```

2. To install the scheduler, execute the following Helm commands:

```
helm repo add yunikorn https://apache.github.io/yunikorn-release
```

```
helm repo update
```

```
helm install yunikorn yunikorn/yunikorn --namespace yunikorn
```

Run a Spark application with YuniKorn scheduler with the Spark operator

1. If you haven't already, complete the steps in the following sections to get set up:
 - a. [Create your cluster and get set up for YuniKorn](#)
 - b. [Install and set up YuniKorn](#)
 - c. [Setting up the Spark operator for Amazon EMR on EKS](#)
 - d. [Install the Spark operator](#)

Include the following arguments when you run the `helm install spark-operator-demo` command:

```
--set batchScheduler.enable=true
--set webhook.enable=true
```

2. Create a SparkApplication definition file spark-pi.yaml.

To use YuniKorn as a scheduler for your jobs, you must add certain annotations and labels to your application definition. The annotations and labels specify the queue for your job and the scheduling strategy that you want to use.

In the following example, the annotation `schedulingPolicyParameters` sets up gang scheduling for the application. Then, the example creates **task groups**, or "gangs" of tasks, to specify the minimum capacity that must be available before scheduling the pods to start the job execution. And finally, it specifies in the task group definition to use node groups with the `"app": "spark"` label, as defined in the [Create your cluster and get set up for YuniKorn](#) section.

```
apiVersion: "sparkoperator.k8s.io/v1beta2"
kind: SparkApplication
metadata:
  name: spark-pi
  namespace: spark-job
spec:
  type: Scala
  mode: cluster
  image: "895885662937.dkr.ecr.us-west-2.amazonaws.com/spark/emr-6.10.0:latest"
  imagePullPolicy: Always
  mainClass: org.apache.spark.examples.SparkPi
  mainApplicationFile: "local:///usr/lib/spark/examples/jars/spark-examples.jar"
  sparkVersion: "3.3.1"
  restartPolicy:
    type: Never
  volumes:
    - name: "test-volume"
      hostPath:
        path: "/tmp"
        type: Directory
  driver:
    cores: 1
    coreLimit: "1200m"
    memory: "512m"
    labels:
```

```

    version: 3.3.1
  annotations:
    yunikorn.apache.org/schedulingPolicyParameters: "placeholderTimeoutSeconds=30
gangSchedulingStyle=Hard"
    yunikorn.apache.org/task-group-name: "spark-driver"
    yunikorn.apache.org/task-groups: |-
      [{
        "name": "spark-driver",
        "minMember": 1,
        "minResource": {
          "cpu": "1200m",
          "memory": "1Gi"
        },
        "nodeSelector": {
          "app": "spark"
        }
      },
      {
        "name": "spark-executor",
        "minMember": 1,
        "minResource": {
          "cpu": "1200m",
          "memory": "1Gi"
        },
        "nodeSelector": {
          "app": "spark"
        }
      }
    ]
  serviceAccount: spark-sa
  volumeMounts:
    - name: "test-volume"
      mountPath: "/tmp"
  executor:
    cores: 1
    instances: 1
    memory: "512m"
    labels:
      version: 3.3.1
    annotations:
      yunikorn.apache.org/task-group-name: "spark-executor"
    volumeMounts:
      - name: "test-volume"
        mountPath: "/tmp"

```


3. Submit the Spark application with the following command. This also creates a SparkApplication object called spark-pi:

```
kubectl apply -f spark-pi.yaml
```

4. Check events for the SparkApplication object with the following command:

```
kubectl describe sparkapplication spark-pi --namespace spark-job
```

The first pod event will show that YuniKorn has scheduled the pods:

Type	Reason	Age	From	Message
Normal	Scheduling	3m12s	yunikorn	spark-operator/org-apache-spark-examples-sparkpi-2a777a88b98b8a95-driver is queued and waiting for allocation
Normal	GangScheduling	3m12s	yunikorn	Pod belongs to the taskGroup spark-driver, it will be scheduled as a gang member
Normal	Scheduled	3m10s	yunikorn	Successfully assigned spark
Normal	PodBindSuccessful	3m10s	yunikorn	Pod spark-operator/
Normal	TaskCompleted	2m3s	yunikorn	Task spark-operator/
Normal	Pulling	3m10s	kubelet	Pulling

Run a Spark application with YuniKorn scheduler with spark-submit

1. First, complete the steps in the [Setting up spark-submit for Amazon EMR on EKS](#) section.
2. Set the values for the following environment variables:

```
export SPARK_HOME=spark-home
export MASTER_URL=k8s://Amazon-EKS-cluster-endpoint
```

3. Submit the Spark application with the following command:

In the following example, the annotation schedulingPolicyParameters sets up gang scheduling for the application. Then, the example creates **task groups**, or "gangs" of tasks, to specify the minimum capacity that must be available before scheduling the pods to start the job execution. And finally, it specifies in the task group definition to use node groups with the "app": "spark" label, as defined in the [Create your cluster and get set up for YuniKorn](#) section.

```

$SPARK_HOME/bin/spark-submit \
  --class org.apache.spark.examples.SparkPi \
  --master $MASTER_URL \
  --conf spark.kubernetes.container.image=895885662937.dkr.ecr.us-
west-2.amazonaws.com/spark/emr-6.10.0:latest \
  --conf spark.kubernetes.authenticate.driver.serviceAccountName=spark-sa \
  --deploy-mode cluster \
  --conf spark.kubernetes.namespace=spark-job \
  --conf spark.kubernetes.scheduler.name=yunikorn \
  --conf spark.kubernetes.driver.annotation.yunikorn.apache.org/
schedulingPolicyParameters="placeholderTimeoutSeconds=30 gangSchedulingStyle=Hard"
\
  --conf spark.kubernetes.driver.annotation.yunikorn.apache.org/task-group-
name="spark-driver" \
  --conf spark.kubernetes.executor.annotation.yunikorn.apache.org/task-group-
name="spark-executor" \
  --conf spark.kubernetes.driver.annotation.yunikorn.apache.org/task-groups='[{
    "name": "spark-driver",
    "minMember": 1,
    "minResource": {
      "cpu": "1200m",
      "memory": "1Gi"
    },
    "nodeSelector": {
      "app": "spark"
    }
  },
  {
    "name": "spark-executor",
    "minMember": 1,
    "minResource": {
      "cpu": "1200m",
      "memory": "1Gi"
    },
    "nodeSelector": {
      "app": "spark"
    }
  }
]'] \
local:///usr/lib/spark/examples/jars/spark-examples.jar 20

```

4. Check events for the SparkApplication object with the following command:

```
kubectl describe pod spark-driver-pod --namespace spark-job
```

The first pod event will show that YuniKorn has scheduled the pods:

Type	Reason	Age	From	Message
----	-----	----	----	-----
Normal	Scheduling	3m12s	yunikorn	spark-operator/org-apache-spark-examples-sparkpi-2a777a88b98b8a95-driver is queued and waiting for allocation
Normal	GangScheduling	3m12s	yunikorn	Pod belongs to the taskGroup spark-driver, it will be scheduled as a gang member
Normal	Scheduled	3m10s	yunikorn	Successfully assigned spark
Normal	PodBindSuccessful	3m10s	yunikorn	Pod spark-operator/
Normal	TaskCompleted	2m3s	yunikorn	Task spark-operator/
Normal	Pulling	3m10s	kubelet	Pulling

Security in Amazon EMR on EKS

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 Amazon EMR, 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 Amazon EMR on EKS. The following topics show you how to configure Amazon EMR on EKS to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Amazon EMR on EKS resources.

Topics

- [Amazon EMR on EKS security best practices](#)
- [Data protection](#)
- [Identity and Access Management](#)
- [Using Amazon EMR on EKS with AWS Lake Formation for fine-grained access control](#)
- [Logging and monitoring](#)
- [Using Amazon S3 Access Grants with Amazon EMR on EKS](#)
- [Compliance validation for Amazon EMR on EKS](#)
- [Resilience in Amazon EMR on EKS](#)
- [Infrastructure security in Amazon EMR on EKS](#)
- [Configuration and vulnerability analysis](#)

- [Connect to Amazon EMR on EKS Using an interface VPC endpoint](#)
- [Set up cross-account access for Amazon EMR on EKS](#)

Amazon EMR on EKS security best practices

Amazon EMR on EKS 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 security best practices, see [Amazon EMR on EKS security best practices](#).

Apply principle of least privilege

Amazon EMR on EKS provides a granular access policy for applications using IAM roles, such as execution roles. These execution roles are mapped to Kubernetes service accounts through the IAM role's trust policy. Amazon EMR on EKS creates pods within a registered Amazon EKS namespace that execute user-provided application code. The job pods running the application code assume the execution role when connecting to other AWS services. We recommend that execution roles be granted only the minimum set of privileges required by the job, such as covering your application and access to log destination. We also recommend auditing the jobs for permissions on a regular basis and upon any change to application code.

Access control list for endpoints

Managed endpoints can be created only for those EKS clusters that have been configured to use at least one private subnet in your VPC. This configuration restricts access to the load balancers created by managed endpoints so that they can only be accessed from your VPC. To further enhance security, we recommend that you configure security groups with these load balancers so that they can restrict incoming traffic to a selected set of IP addresses.

Get the latest security updates for custom images

To use custom images with Amazon EMR on EKS, you can install any binaries and libraries on the image. You are responsible for the security patching of the binaries you add to the image. Amazon

EMR on EKS images are regularly patched with latest security patches. To get the latest image, you must rebuild the custom images whenever there is a new base image version of the Amazon EMR release. For more information, see [Amazon EMR on EKS releases](#) and [Details for selecting a base image URI](#).

Limit pod credential access

Kubernetes supports several methods of assigning credentials to a pod. Provisioning multiple credentials providers can increase the complexity of your security model. Amazon EMR on EKS has adopted the use of [IAM roles for services accounts \(IRSA\)](#) as a standard credential provider within a registered EKS namespace. Other methods are not supported, including [kube2iam](#), [kiam](#) and using an EC2 instance profile of the instance running on the cluster.

Isolate untrusted application code

Amazon EMR on EKS does not inspect the integrity of the application code submitted by users of the system. If you are running a multi-tenanted virtual cluster that is configured using multiple execution roles that can be used to submit jobs by untrusted tenants running arbitrary code, there is a risk of a malicious application escalating its privileges. In this situation, consider isolating execution roles with similar privileges into a different virtual cluster.

Role-based access control (RBAC) permissions

Administrators should strictly control Role-based access control (RBAC) permissions for Amazon EMR on EKS managed namespaces. At a minimum, the following permissions should not be granted to job submitters in Amazon EMR on EKS managed namespaces.

- Kubernetes RBAC permissions to modify configmap - because Amazon EMR on EKS uses Kubernetes configmaps to generate managed pod templates that have the managed service account name. This attribute should not be mutated.
- Kubernetes RBAC permissions to exec into Amazon EMR on EKS pods - to avoid giving access to managed pod templates that have the managed SA name. This attribute should not be mutated. This permission can also give access to the JWT token mounted into the pod which can then be used to retrieve the execution role credentials.
- Kubernetes RBAC permissions to create pods - to prevent users from creating pods using a Kubernetes ServiceAccount which may be mapped to an IAM role with more AWS privileges than the user.

- Kubernetes RBAC permissions to deploy mutating webhook - to prevent users from using the mutating webhook to mutate Kubernetes ServiceAccount name for pods created by Amazon EMR on EKS.
- Kubernetes RBAC permissions to read Kubernetes secrets - to prevent users from reading confidential data stored in these secrets.

Restrict access to nodegroup IAM role or instance profile credentials

- We recommend that you assign minimum AWS permissions to nodegroup's IAM role(s). This helps to avoid privilege escalation by code that may run using instance profile credentials of EKS worker nodes.
- To completely block access to instance profile credentials to all pods that runs in Amazon EMR on EKS managed namespaces, we recommend that you run `iptables` commands on EKS nodes. For more information, see [Restricting access to Amazon EC2 instance profile credentials](#). However, it is important to properly scope your service account IAM roles so that your pods have all of the necessary permissions. For example, the node IAM role is assigned permissions to pull container images from Amazon ECR. If a pod isn't assigned those permissions, the pod can't pull container images from Amazon ECR. The VPC CNI plugin also needs to be updated. For more information, see [Walkthrough: Updating the VPC CNI plugin to use IAM roles for service accounts](#).

Data protection

The AWS [shared responsibility model](#) applies to data protection in Amazon EMR on EKS. 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. This content includes 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 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 recommend TLS 1.2 or later.
- 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 S3.
- Use Amazon EMR on EKS encryption options to encrypt data at rest and in transit.
- 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 Amazon EMR on EKS or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into Amazon EMR on EKS 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.

Encryption at rest

Data encryption helps prevent unauthorized users from reading data on a cluster and associated data storage systems. This includes data saved to persistent media, known as data at rest, and data that may be intercepted as it travels the network, known as data in transit.

Data encryption requires keys and certificates. You can choose from several options, including keys managed by AWS Key Management Service, keys managed by Amazon S3, and keys and certificates from custom providers that you supply. When using AWS KMS as your key provider, charges apply for the storage and use of encryption keys. For more information, see [AWS KMS Pricing](#).

Before you specify encryption options, decide on the key and certificate management systems you want to use. Then create the keys and certificates for the custom providers that you specify as part of encryption settings.

Encryption at rest for EMRFS data in Amazon S3

Amazon S3 encryption works with EMR File System (EMRFS) objects read from and written to Amazon S3. You specify Amazon S3 server-side encryption (SSE) or client-side encryption (CSE)

as the **Default encryption mode** when you enable encryption at rest. Optionally, you can specify different encryption methods for individual buckets using **Per bucket encryption overrides**. Regardless of whether Amazon S3 encryption is enabled, Transport Layer Security (TLS) encrypts the EMRFS objects in transit between EMR cluster nodes and Amazon S3. For in-depth information about Amazon S3 encryption, see [Protecting Data Using Encryption](#) in the Amazon Simple Storage Service Developer Guide.

Note

When you use AWS KMS, charges apply for the storage and use of encryption keys. For more information, see [AWS KMS Pricing](#).

Amazon S3 server-side encryption

When you set up Amazon S3 server-side encryption, Amazon S3 encrypts data at the object level as it writes the data to disk and decrypts the data when it is accessed. For more information about SSE, see [Protecting Data Using Server-Side Encryption](#) in the Amazon Simple Storage Service Developer Guide.

You can choose between two different key management systems when you specify SSE in Amazon EMR on EKS:

- **SSE-S3** - Amazon S3 manages keys for you.
- **SSE-KMS** - You use an AWS KMS key to set up with policies suitable for Amazon EMR on EKS.

SSE with customer-provided keys (SSE-C) is not available for use with Amazon EMR on EKS.

Tip

To reduce AWS KMS costs when using SSE-KMS, consider enabling Amazon S3 Bucket Keys on your Amazon S3 buckets. Amazon S3 Bucket Keys use a short-lived bucket-level key to reduce AWS KMS API calls by up to 99 percent. Before enabling Amazon S3 Bucket Keys, review your IAM and AWS KMS key policies – the encryption context changes from the Amazon S3 object ARN to the bucket ARN, which may affect policies that use the object ARN for access control. For more information, see [Reducing the cost of SSE-KMS with Amazon S3 Bucket Keys](#) in the *Amazon Simple Storage Service User Guide*.

Amazon S3 client-side encryption

With Amazon S3 client-side encryption, the Amazon S3 encryption and decryption takes place in the EMRFS client on your cluster. Objects are encrypted before being uploaded to Amazon S3 and decrypted after they are downloaded. The provider you specify supplies the encryption key that the client uses. The client can use keys provided by AWS KMS (CSE-KMS) or a custom Java class that provides the client-side root key (CSE-C). The encryption specifics are slightly different between CSE-KMS and CSE-C, depending on the specified provider and the metadata of the object being decrypted or encrypted. For more information about these differences, see [Protecting Data Using Client-Side Encryption](#) in the Amazon Simple Storage Service Developer Guide.

Note

Amazon S3 CSE only ensures that EMRFS data exchanged with Amazon S3 is encrypted; not all data on cluster instance volumes is encrypted. Furthermore, because Hue does not use EMRFS, objects that the Hue S3 File Browser writes to Amazon S3 are not encrypted.

Local disk encryption

Apache Spark supports encrypting temporary data written to local disks. This covers shuffle files, shuffle spills, and data blocks stored on disk for both caching and broadcast variables. It does not cover encrypting output data generated by applications with APIs such as `saveAsHadoopFile` or `saveAsTable`. It also may not cover temporary files created explicitly by the user. For more information, see [Local Storage Encryption](#) in the Spark documentation. Spark does not support encrypted data on local disk, such as intermediate data written to a local disk by an executor process when the data does not fit in memory. Data that is persisted to disk is scoped to the job runtime, and the key that is used to encrypt the data is generated dynamically by Spark for every job run. Once the Spark job terminates, no other process can decrypt the data.

For driver and executor pod, you encrypt data at rest that is persisted to the mounted volume. There are three different AWS native storage options you can use with Kubernetes: [EBS](#), [EFS](#), and [FSx for Lustre](#). All three offer encryption at rest using a service managed key or an AWS KMS key. For more information see the [EKS Best Practices Guide](#). With this approach, all data persisted to the mounted volume is encrypted.

Key management

You can configure KMS to automatically rotate your KMS keys. This rotates your keys once a year while saving old keys indefinitely so that your data can still be decrypted. For additional information, see [Rotating AWS KMS keys](#).

Encryption in transit

Several encryption mechanisms are enabled with in-transit encryption. These are open-source features, are application-specific, and may vary by Amazon EMR on EKS release. The following application-specific encryption features can be enabled with Amazon EMR on EKS:

- Spark
 - Internal RPC communication between Spark components, such as the block transfer service and the external shuffle service, is encrypted using the AES-256 cipher in Amazon EMR versions 5.9.0 and later. In earlier releases, internal RPC communication is encrypted using SASL with DIGEST-MD5 as the cipher.
 - HTTP protocol communication with user interfaces such as Spark History Server and HTTPS-enabled file servers is encrypted using Spark's SSL configuration. For more information, see [SSL Configuration](#) in Spark documentation.

For more information, see [Spark security settings](#).

- You should allow only encrypted connections over HTTPS (TLS) using [the aws:SecureTransport condition](#) on Amazon S3 bucket IAM policies.
- Query results that stream to JDBC or ODBC clients are encrypted using TLS.

Identity and Access Management

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 Amazon EMR on EKS resources. IAM is an AWS service that you can use with no additional charge.

Topics

- [Audience](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)

- [How Amazon EMR on EKS works with IAM](#)
- [Using service-linked roles for Amazon EMR on EKS](#)
- [Managed policies for Amazon EMR on EKS](#)
- [Using job execution roles with Amazon EMR on EKS](#)
- [Identity-based policy examples for Amazon EMR on EKS](#)
- [Policies for tag-based access control](#)
- [Troubleshooting Amazon EMR on EKS identity and access](#)

Audience

How you use AWS Identity and Access Management (IAM) differs based on your role:

- **Service user** - request permissions from your administrator if you cannot access features (see [Troubleshooting Amazon EMR on EKS identity and access](#))
- **Service administrator** - determine user access and submit permission requests (see [How Amazon EMR on EKS works with IAM](#))
- **IAM administrator** - write policies to manage access (see [Identity-based policy examples for Amazon EMR on EKS](#))

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be authenticated as the AWS account root user, an IAM user, or by assuming an IAM role.

You can sign in as a federated identity using credentials from an identity source like AWS IAM Identity Center (IAM Identity Center), single sign-on authentication, or Google/Facebook credentials. For more information about signing in, see [How to sign in to your AWS account](#) in the *AWS Sign-In User Guide*.

For programmatic access, AWS provides an SDK and CLI to cryptographically sign requests. For more information, see [AWS Signature Version 4 for API requests](#) in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity called the AWS account *root user* that has complete access to all AWS services and resources. We strongly recommend that you

don't use the root user for everyday tasks. For tasks that require root user credentials, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

Federated identity

As a best practice, require human users to use federation with an identity provider to access AWS services using temporary credentials.

A *federated identity* is a user from your enterprise directory, web identity provider, or Directory Service that accesses AWS services using credentials from an identity source. Federated identities assume roles that provide temporary credentials.

For centralized access management, we recommend AWS IAM Identity Center. For more information, see [What is IAM Identity Center?](#) in the *AWS IAM Identity Center User Guide*.

IAM users and groups

An [IAM user](#) is an identity with specific permissions for a single person or application. We recommend using temporary credentials instead of IAM users with long-term credentials. For more information, see [Require human users to use federation with an identity provider to access AWS using temporary credentials](#) in the *IAM User Guide*.

An [IAM group](#) specifies a collection of IAM users and makes permissions easier to manage for large sets of users. For more information, see [Use cases for IAM users](#) in the *IAM User Guide*.

IAM roles

An [IAM role](#) is an identity with specific permissions that provides temporary credentials. You can assume a role by [switching from a user to an IAM role \(console\)](#) or by calling an AWS CLI or AWS API operation. For more information, see [Methods to assume a role](#) in the *IAM User Guide*.

IAM roles are useful for federated user access, temporary IAM user permissions, cross-account access, cross-service access, and applications running on Amazon EC2. For more information, see [Cross account resource access in IAM](#) 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 defines permissions when associated with an identity or resource. AWS evaluates these policies when a principal makes a request. Most policies are stored in AWS as JSON documents. For

more information about JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Using policies, administrators specify who has access to what by defining which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. An IAM administrator creates IAM policies and adds them to roles, which users can then assume. IAM policies define permissions regardless of the method used to perform the operation.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you attach to an identity (user, group, or role). These policies control what actions identities can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Define custom IAM permissions with customer managed policies](#) in the *IAM User Guide*.

Identity-based policies can be *inline policies* (embedded directly into a single identity) or *managed policies* (standalone policies attached to multiple identities). To learn how to choose between managed and inline policies, see [Choose 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 include 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. You must [specify a principal](#) in a resource-based policy.

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.

Other policy types

AWS supports additional policy types that can set the maximum permissions granted by more common policy types:

- **Permissions boundaries** – Set the maximum permissions that an identity-based policy can grant to an IAM entity. For more information, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.

- **Service control policies (SCPs)** – Specify the maximum permissions for an organization or organizational unit in AWS Organizations. For more information, see [Service control policies](#) in the *AWS Organizations User Guide*.
- **Resource control policies (RCPs)** – Set the maximum available permissions for resources in your accounts. For more information, see [Resource control policies \(RCPs\)](#) in the *AWS Organizations User Guide*.
- **Session policies** – Advanced policies passed as a parameter when creating a temporary session for a role or federated user. 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 Amazon EMR on EKS works with IAM

Before you use IAM to manage access to Amazon EMR on EKS, learn what IAM features are available to use with Amazon EMR on EKS.

IAM features you can use with Amazon EMR on EKS

IAM feature	Amazon EMR on EKS support
Identity-based policies	Yes
Resource-based policies	No
Policy actions	Yes
Policy resources	Yes
Policy condition keys	Yes
ACLs	No
ABAC (tags in policies)	Yes
Temporary credentials	Yes

IAM feature	Amazon EMR on EKS support
Principal permissions	Yes
Service roles	No
Service-linked roles	Yes

To get a high-level view of how Amazon EMR on EKS 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 Amazon EMR on EKS

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 [Define custom IAM permissions with customer managed 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. 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 Amazon EMR on EKS

To view examples of Amazon EMR on EKS identity-based policies, see [Identity-based policy examples for Amazon EMR on EKS](#).

Resource-based policies within Amazon EMR on EKS

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. For more information, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Policy actions for Amazon EMR on EKS

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. Include actions in a policy to grant permissions to perform the associated operation.

To see a list of Amazon EMR on EKS actions, see [Actions, resources, and condition keys for Amazon EMR on EKS](#) in the *Service Authorization Reference*.

Policy actions in Amazon EMR on EKS use the following prefix before the action:

```
emr-containers
```

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [  
    "emr-containers:action1",  
    "emr-containers:action2"  
]
```

To view examples of Amazon EMR on EKS identity-based policies, see [Identity-based policy examples for Amazon EMR on EKS](#).

Policy resources for Amazon EMR on EKS

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. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). For actions that don't support resource-level permissions, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*" 
```

To see a list of Amazon EMR on EKS resource types and their ARNs, see [Resources defined by Amazon EMR on EKS](#) in the *Service Authorization Reference*. To learn which actions you can specify the ARN of each resource, see [Actions, resources, and condition keys for Amazon EMR on EKS](#).

To view examples of Amazon EMR on EKS identity-based policies, see [Identity-based policy examples for Amazon EMR on EKS](#).

Policy condition keys for Amazon EMR on EKS

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 specifies when statements execute based on defined criteria. 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. To see all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*.

To see a list of Amazon EMR on EKS condition keys and to learn which actions and resources you can use a condition key, see [Actions, resources, and condition keys for Amazon EMR on EKS](#) in the *Service Authorization Reference*.

To view examples of Amazon EMR on EKS identity-based policies, see [Identity-based policy examples for Amazon EMR on EKS](#).

Access control lists (ACLs) in Amazon EMR on EKS

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.

Attribute-based access control (ABAC) with Amazon EMR on EKS

Supports ABAC (tags in policies)	Yes
----------------------------------	-----

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes called tags. You can attach tags to IAM entities and AWS resources, then design ABAC policies to allow operations when the principal's tag matches the tag on the resource.

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 [Define permissions with ABAC authorization](#) 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 Amazon EMR on EKS

Supports temporary credentials: Yes

Temporary credentials provide short-term access to AWS resources and are automatically created when you use federation or switch roles. AWS recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see [Temporary security credentials in IAM](#) and [AWS services that work with IAM](#) in the *IAM User Guide*.

Cross-service principal permissions for Amazon EMR on EKS

Supports forward access sessions (FAS): Yes

Forward access sessions (FAS) use the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. For policy details when making FAS requests, see [Forward access sessions](#).

Service roles for Amazon EMR on EKS

Supports service roles	No
------------------------	----

Service-linked roles for Amazon EMR on EKS

Supports service-linked roles	Yes
-------------------------------	-----

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.

Using service-linked roles for Amazon EMR on EKS

Amazon EMR on EKS uses AWS Identity and Access Management (IAM) [service-linked roles](#). A service-linked role is a unique type of IAM role that is linked directly to Amazon EMR on EKS. Service-linked roles are predefined by Amazon EMR on EKS and include all the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes setting up Amazon EMR on EKS easier because you don't have to manually add the necessary permissions. Amazon EMR on EKS defines the permissions of its service-linked roles, and unless defined otherwise, only Amazon EMR on EKS can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

You can delete a service-linked role only after first deleting their related resources. This protects your Amazon EMR on EKS resources because you can't inadvertently remove permission to access the resources.

For information about other services that support service-linked roles, see [AWS Services That Work with IAM](#) and look for the services that have **Yes** in the **Service-Linked Role** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

Service-linked role permissions for Amazon EMR on EKS

Amazon EMR on EKS uses the service-linked role named **AWSServiceRoleForAmazonEMRContainers**.

The `AWSServiceRoleForAmazonEMRContainers` service-linked role trusts the following services to assume the role:

- `emr-containers.amazonaws.com`

The role permissions policy `AmazonEMRContainersServiceRolePolicy` allows Amazon EMR on EKS to complete a set of actions on the specified resources, as the following policy statement demonstrates.

 Note

Managed policy contents change, so the policy shown here may be out-of-date. View the most up-to-date policy documentation at [AmazonEMRContainersServiceRolePolicy](#) in the *AWS Managed Policy Reference Guide*.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "eks:DescribeCluster",
        "eks:ListNodeGroups",
        "eks:DescribeNodeGroup",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSubnets",
        "ec2:DescribeSecurityGroups",
        "elasticloadbalancing:DescribeInstanceHealth",
        "elasticloadbalancing:DescribeLoadBalancers",
        "elasticloadbalancing:DescribeTargetGroups",
        "elasticloadbalancing:DescribeTargetHealth",
        "eks:ListPodIdentityAssociations",
        "eks:DescribePodIdentityAssociation"
      ],
      "Resource": [
        "*"
      ],
    },
  ],
}
```

```

    "Sid": "AllowEKSDescribecluster"
  },
  {
    "Effect": "Allow",
    "Action": [
      "acm:ImportCertificate",
      "acm:AddTagsToCertificate"
    ],
    "Resource": [
      "*"
    ],
    "Condition": {
      "StringEquals": {
        "aws:RequestTag/emr-container:endpoint:managed-certificate": "true"
      }
    },
    "Sid": "AllowACMImportcertificate"
  },
  {
    "Effect": "Allow",
    "Action": [
      "acm:DeleteCertificate"
    ],
    "Resource": [
      "*"
    ],
    "Condition": {
      "StringEquals": {
        "aws:ResourceTag/emr-container:endpoint:managed-certificate": "true"
      }
    },
    "Sid": "AllowACMDeletecertificate"
  }
]
}

```

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see [Service-Linked Role Permissions](#) in the *IAM User Guide*.

Creating a service-linked role for Amazon EMR on EKS

You don't need to manually create a service-linked role. When you create a virtual cluster, Amazon EMR on EKS creates the service-linked role for you.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you create a virtual cluster, Amazon EMR on EKS creates the service-linked role for you again.

You can also use the IAM console to create a service-linked role with the **Amazon EMR on EKS** use case. In the AWS CLI or the AWS API, create a service-linked role with the `emr-containers.amazonaws.com` service name. For more information, see [Creating a Service-Linked Role](#) in the *IAM User Guide*. If you delete this service-linked role, you can use this same process to create the role again.

Editing a service-linked role for Amazon EMR on EKS

Amazon EMR on EKS does not allow you to edit the `AWSServiceRoleForAmazonEMRContainers` service-linked role. After you create a service-linked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see [Editing a Service-Linked Role](#) in the *IAM User Guide*.

Deleting a service-linked role for Amazon EMR on EKS

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don't have an unused entity that is not actively monitored or maintained. However, you must clean up the resources for your service-linked role before you can manually delete it.

Note

If the Amazon EMR on EKS service is using the role when you try to delete the resources, then the deletion might fail. If that happens, wait for a few minutes and try the operation again.

To delete Amazon EMR on EKS resources used by the `AWSServiceRoleForAmazonEMRContainers`

1. Open the Amazon EMR console.
2. Choose a virtual cluster.
3. On the `Virtual Cluster` page choose **Delete**.
4. Repeat this procedure for any other virtual clusters in your account.

To manually delete the service-linked role using IAM

Use the IAM console, the AWS CLI, or the AWS API to delete the `AWSServiceRoleForAmazonEMRContainers` service-linked role. For more information, see [Deleting a Service-Linked Role](#) in the *IAM User Guide*.

Supported Regions for Amazon EMR on EKS service-linked roles

Amazon EMR on EKS supports using service-linked roles in all of the Regions where the service is available. For more information, see [Amazon EMR on EKS service endpoints and quotas](#).

Managed policies for Amazon EMR on EKS

View details about updates to AWS managed policies for Amazon EMR on EKS since March 1, 2021.

Change	Description	Date
<code>AmazonEMRContainersServiceRolePolicy</code> - Added read permissions to list EKS pod identity associations in a cluster, and another read permission to return descriptive information about pod identity associations in a cluster. For more information, see	The following permissions are added to the policy: <code>eks:ListPodIdentityAssociations</code>, <code>eks:DescribePodIdentityAssociation</code>.	February 3, 2023

Change	Description	Date
AmazonEMRContainerServiceRolePolicy .		
AmazonEMRContainerServiceRolePolicy - Added permissions to describe and list Amazon EKS nodegroups, describe load balancer target groups, and describe load balancer target health.	The following permissions are added to the policy: <code>eks:ListNodeGroups</code> , <code>eks:DescribeNodeGroup</code> , <code>elasticloadbalancing:DescribeTargetGroups</code> , <code>elasticloadbalancing:DescribeTargetHealth</code> .	March 13, 2023
AmazonEMRContainerServiceRolePolicy - Added permissions to import and delete certificates in AWS Certificate Manager.	The following permissions are added to the policy: <code>acm:ImportCertificate</code> , <code>acm:AddTagsToCertificate</code> , <code>acm:DeleteCertificate</code> .	Dec 3, 2021
Amazon EMR on EKS started tracking changes	Amazon EMR on EKS started tracking changes for its AWS managed policies.	March 1, 2021

Using job execution roles with Amazon EMR on EKS

To use the `StartJobRun` command to submit a job run on an EKS cluster, you must first onboard a job execution role to be used with a virtual cluster. For more information, see [Create a job execution role](#) in [Setting up Amazon EMR on EKS](#). You can also follow the instructions in the [Create IAM Role for job execution](#) section of the Amazon EMR on EKS Workshop.

The following permissions must be included in the trust policy for the job execution role.

JSON

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

```

{
  "Effect": "Allow",
  "Principal": {
    "Federated": "arn:aws:iam::AWS_ACCOUNT_ID:oidc-provider/OIDC_PROVIDER"
  },
  "Action": "sts:AssumeRoleWithWebIdentity",
  "Condition": {
    "StringEquals": {
      "OIDC_PROVIDER:aud": "sts.amazonaws.com"
    },
    "StringLike": {
      "OIDC_PROVIDER:sub": "system:serviceaccount:NAMESPACE:emr-containers-
sa-*-*AWS_ACCOUNT_ID-BASE36_ENCODED_ROLE_NAME"
    }
  }
}

```

The trust policy in the preceding example grants permissions only to an Amazon EMR managed Kubernetes service account with a name that matches the `emr-containers-sa-*-*AWS_ACCOUNT_ID-BASE36_ENCODED_ROLE_NAME` pattern. Service accounts with this pattern will be automatically created at job submission, and scoped to the namespace where you submit the job. This trust policy allows these service accounts to assume the execution role and get the temporary credentials of the execution role. Service accounts from a different Amazon EKS cluster or from a different namespace within the same EKS cluster are restricted from assuming the execution role.

You can run the following command to automatically update the trust policy in the format given above.

```

aws emr-containers update-role-trust-policy \
  --cluster-name cluster \
  --namespace namespace \
  --role-name iam_role_name_for_job_execution

```

Controlling access to the execution role

An administrator for your Amazon EKS cluster can create a multi-tenant Amazon EMR on EKS virtual cluster to which an IAM administrator can add multiple execution roles. Because untrusted

tenants can use these execution roles to submit jobs that run arbitrary code, you might want to restrict those tenants so that they can't run code that gains the permissions assigned to one or more of these execution roles. To restrict the IAM policy attached to an IAM identity, the IAM administrator can use the optional Amazon Resource Name (ARN) condition key `emr-containers:ExecutionRoleArn`. This condition accepts a list of execution role ARNs that have permissions to the virtual cluster, as the following permissions policy demonstrates.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "emr-containers:StartJobRun"
      ],
      "Resource": [
        "arn:aws:emr-containers:*:*:/virtualclusters/VIRTUAL_CLUSTER_ID"
      ],
      "Condition": {
        "ArnEquals": {
          "emr-containers:ExecutionRoleArn": [
            "arn:aws:iam:*:*:role/execution_role_name_1",
            "arn:aws:iam:*:*:role/execution_role_name_2"
          ]
        }
      },
      "Sid": "AllowEMRCONTAINERSStartjobrun"
    }
  ]
}
```

If you want to allow all execution roles that begin with a particular prefix, such as `MyRole`, you can replace the condition operator `ArnEquals` with the `ArnLike` operator, and you can replace the `execution_role_arn` value in the condition with a wildcard `*` character. For example, `arn:aws:iam:*:AWS_ACCOUNT_ID:role/MyRole*`. All other [ARN condition keys](#) are also supported.

Note

With Amazon EMR on EKS, you can't grant permissions to execution roles based on tags or attributes. Amazon EMR on EKS doesn't support tag-based access control (TBAC) or attribute-based access control (ABAC) for execution roles.

Identity-based policy examples for Amazon EMR on EKS

By default, users and roles don't have permission to create or modify Amazon EMR on EKS resources. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see [Create IAM policies \(console\)](#) in the *IAM User Guide*.

For details about actions and resource types defined by Amazon EMR on EKS, including the format of the ARNs for each of the resource types, see [Actions, resources, and condition keys for Amazon EMR on EKS](#) in the *Service Authorization Reference*.

Topics

- [Policy best practices](#)
- [Using the Amazon EMR on EKS console](#)
- [Allow users to view their own permissions](#)

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Amazon EMR on EKS 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 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 [Validate policies with IAM Access Analyzer](#) 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 [Secure API access with MFA](#) in the *IAM User Guide*.

For more information about best practices in IAM, see [Security best practices in IAM](#) in the *IAM User Guide*.

Using the Amazon EMR on EKS console

To access the Amazon EMR on EKS console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Amazon EMR on EKS 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 Amazon EMR on EKS console, also attach the Amazon EMR on EKS ConsoleAccess or ReadOnly AWS managed policy to the entities. For more information, see [Adding permissions to a user](#) in the *IAM User Guide*.

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewOwnUserInfo",
      "Effect": "Allow",
      "Action": [
        "iam:GetUserPolicy",
        "iam:ListGroupsForUser",
        "iam:ListAttachedUserPolicies",
        "iam:ListUserPolicies",
        "iam:GetUser"
      ],
      "Resource": ["arn:aws:iam::*:user/${aws:username}"]
    },
    {
      "Sid": "NavigateInConsole",
      "Effect": "Allow",
      "Action": [
        "iam:GetGroupPolicy",
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedGroupPolicies",
        "iam:ListGroupPolicies",
        "iam:ListPolicyVersions",
        "iam:ListPolicies",
        "iam:ListUsers"
      ],
      "Resource": "*"
    }
  ]
}
```

Policies for tag-based access control

You can use conditions in your identity-based policy to control access to virtual clusters and job runs based on tags. For more information about tagging, see [Tagging your Amazon EMR on EKS resources](#).

The following examples demonstrate different scenarios and ways to use condition operators with Amazon EMR on EKS condition keys. These IAM policy statements are intended for demonstration purposes only and should not be used in production environments. There are multiple ways to combine policy statements to grant and deny permissions according to your requirements. For more information about planning and testing IAM policies, see the [IAM user Guide](#).

Important

Explicitly denying permission for tagging actions is an important consideration. This prevents users from tagging a resource and thereby granting themselves permissions that you did not intend to grant. If tagging actions for a resource are not denied, a user can modify tags and circumvent the intention of the tag-based policies. For an example of a policy that denies tagging actions, see [Deny access to add and remove tags](#).

The examples below demonstrate identity-based permissions policies that are used to control the actions that are allowed with Amazon EMR on EKS virtual clusters.

Allow actions only on resources with specific tag values

In the following policy example, the StringEquals condition operator tries to match dev with the value for the tag department. If the tag department hasn't been added to the virtual cluster, or doesn't contain the value dev, the policy doesn't apply, and the actions aren't allowed by this policy. If no other policy statements allow the actions, the user can only work with virtual clusters that have this tag with this value.

JSON

```
{
```

```
"Version": "2012-10-17",
"Statement": [
  {
    "Effect": "Allow",
    "Action": [
      "emr-containers:DescribeVirtualCluster"
    ],
    "Resource": [
      "*"
    ],
    "Condition": {
      "StringEquals": {
        "aws:ResourceTag/department": "dev"
      }
    },
    "Sid": "AllowEMRCONTAINERDescribevirtualcluster"
  }
]
```

You can also specify multiple tag values using a condition operator. For example, to allow actions on virtual clusters where the department tag contains the value `dev` or `test`, you could replace the condition block in the earlier example with the following.

```
"Condition": {
  "StringEquals": {
    "aws:ResourceTag/department": ["dev", "test"]
  }
}
```

Require tagging when a resource is created

In the example below, the tag needs to be applied when creating the virtual cluster.

JSON

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



```

    "Action": [
      "emr-containers:CreateVirtualCluster"
    ],
    "Resource": [
      "*"
    ],
    "Condition": {
      "StringEquals": {
        "aws:RequestTag/department": "dev"
      }
    },
    "Sid": "AllowEMRCONTAINERSCreatevirtualcluster"
  }
]
}

```

The following policy statement allows a user to create a virtual cluster only if the cluster has a department tag, which can contain any value.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "emr-containers:CreateVirtualCluster"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "Null": {
          "aws:RequestTag/department": "false"
        }
      },
      "Sid": "AllowEMRCONTAINERSCreatevirtualcluster"
    }
  ]
}

```

Deny access to add and remove tags

The effect of this policy is to deny a user the permission to add or remove any tags on virtual clusters that are tagged with a department tag that contains the dev value.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": [
        "emr-containers:TagResource",
        "emr-containers:UntagResource"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringNotEquals": {
          "aws:ResourceTag/department": "dev"
        }
      },
      "Sid": "AllowEMRCONTAINERSTagresource"
    }
  ]
}
```

Troubleshooting Amazon EMR on EKS identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon EMR on EKS and IAM.

Topics

- [I am not authorized to perform an action in Amazon EMR on EKS](#)
- [I am not authorized to perform iam:PassRole](#)
- [I want to allow people outside of my AWS account to access my Amazon EMR on EKS resources](#)

I am not authorized to perform an action in Amazon EMR on EKS

If the AWS Management Console tells you that you're not authorized to perform an action, then you must contact your administrator for assistance. Your administrator is the person that provided you with your user name and password.

The following example error occurs when the `mateojackson` user tries to use the console to view details about a fictional `my-example-widget` resource but does not have the fictional `emr-containers:GetWidget` permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform: emr-containers:GetWidget on resource: my-example-widget
```

In this case, Mateo asks his administrator to update his policies to allow him to access the `my-example-widget` resource using the `emr-containers:GetWidget` action.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to Amazon EMR on EKS.

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 Amazon EMR on EKS. 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 Amazon EMR on EKS resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Amazon EMR on EKS supports these features, see [How Amazon EMR on EKS 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*.

Using Amazon EMR on EKS with AWS Lake Formation for fine-grained access control

With Amazon EMR release 7.7 and higher, you can leverage AWS Lake Formation to apply fine-grained access controls on AWS Glue Data Catalog tables that are backed by Amazon S3 buckets. This capability lets you configure table, row, column, and cell-level access controls for read and write queries within your Amazon EMR on EKS Spark Jobs.


Topics

- [How Amazon EMR on EKS works with AWS Lake Formation](#)
- [Enable Lake Formation with Amazon EMR on EKS](#)
- [Considerations and limitations](#)
- [Troubleshooting](#)

How Amazon EMR on EKS works with AWS Lake Formation

Using Amazon EMR on EKS with Lake Formation lets you enforce a layer of permissions on each Spark Job to apply Lake Formation permission control when Amazon EMR on EKS executes jobs. Amazon EMR on EKS uses [Spark resource profiles](#) to create two profiles to effectively execute jobs. The User Profile executes user-supplied code, while the system profile enforces Lake Formation policies. Each Lake Formation enabled Job utilizes two Spark drivers, one for the User profile, and another for the System profile. For more information, see [What is AWS Lake Formation](#).

The following is a high-level overview of how Amazon EMR on EKS gets access to data protected by Lake Formation security policies.



The following steps describe this process:

1. A user submits a Spark Job to an AWS Lake Formation-enabled Amazon EMR on EKS virtual cluster.

2. The Amazon EMR on EKS service sets up the User Driver and runs the job in the User Profile. The User Driver runs a lean version of Spark that has no ability to launch tasks, requests executors, access Amazon S3 or the Glue Data Catalog. It only builds a Job plan.
3. The Amazon EMR on EKS service sets up a second driver called a System Driver and runs it in the System Profile (with a privileged identity). Amazon EKS sets up an encrypted TLS channel between the two drivers for communication. The User Driver uses the channel to send the job plans to the System Driver. The System Driver does not run user-submitted code. It runs full Spark and communicates with Amazon S3 and the Data Catalog for data access. It requests executors and compiles the Job Plan into a sequence of execution stages.
4. Amazon EMR on EKS service then runs the stages on executors. User Code in any stage is run exclusively on User profile executors.
5. Stages that read data from Data Catalog tables protected by Lake Formation or those that apply security filters are delegated to System executors.

Enable Lake Formation with Amazon EMR on EKS

With Amazon EMR release 7.7 and higher, you can leverage AWS Lake Formation to apply fine-grained access controls on Data Catalog tables that are backed by Amazon S3. This capability lets you configure table, row, column, and cell level access controls for read and write queries within your Amazon EMR on EKS Spark Jobs.

This section covers how to create a security configuration and set up Lake Formation to work with Amazon EMR. It also describes how to create a virtual cluster with the Security Configuration that you created for Lake Formation. These sections are meant to be completed in sequence.

Step 1: Set up Lake Formation-based column, row, or cell-level permissions

First, to apply row and column level permissions with Lake Formation, the data lake administrator for Lake Formation must set the **LakeFormationAuthorizedCaller** Session Tag. Lake Formation uses this session tag to authorize callers and provide access to the data lake.

Navigate to the AWS Lake Formation console and select the **Application integration settings** option from the **Administration** section in the sidebar. Then, check the box **Allow external engines to filter data in Amazon S3 locations registered with Lake Formation**. Add the **AWS Account IDs** where the Spark Jobs would be running, and the **Session tag Values**.

Application integration settings [Learn more](#)

Application integration settings

Use the options below to control which third-party engines are allowed to read and filter data in Amazon S3 locations registered with Lake Formation.

☒ Allow external engines to filter data in Amazon S3 locations registered with Lake Formation

Check this box to allow third-party engines to access data in Amazon S3 locations that are registered with Lake Formation.

Session tag values

Enter one or more strings that match the LakeFormationAuthorizedCaller session tag defined for third-party engines.

Clear all

EMR on EKS Engine X

Enter one or several string values separated by comma.

AWS account IDs

Enter the external AWS account IDs from where third-party engines are allowed to access locations registered with Lake Formation.

Clear all

012345678901 X
Account

Enter one or more AWS account IDs. Press enter after each ID.

☐ Allow external engines to access data in Amazon S3 locations with full table access

When you enable this option, Lake Formation will return credentials to the integrated application directly without IAM session tag validation.

Cancel

Save

Note that the **LakeFormationAuthorizedCaller** Session Tag passed here is passed in the **SecurityConfiguration** later when you set up IAM roles, in section 3.

Step 2: Setup EKS RBAC permissions

Second, you set up permissions for role-based access control.

Provide EKS Cluster Permissions to the Amazon EMR on EKS service

The Amazon EMR on EKS Service must have EKS Cluster Role permissions so that it can create cross namespace permissions for the System Driver to spin off User executors in the User namespace.

Create Cluster Role

This sample defines permissions for a collection of resources.

```
vim emr-containers-cluster-role.yaml
```

```

---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: emr-containers
rules:
  - apiGroups: [""]
    resources: ["namespaces"]
    verbs: ["get"]
  - apiGroups: [""]
    resources: ["serviceaccounts", "services", "configmaps", "events", "pods", "pods/
log"]
    verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"deletecollection", "annotate", "patch", "label"]
  - apiGroups: [""]
    resources: ["secrets"]
    verbs: ["create", "patch", "delete", "watch"]
  - apiGroups: ["apps"]
    resources: ["statefulsets", "deployments"]
    verbs: ["get", "list", "watch", "describe", "create", "edit", "delete", "annotate",
"patch", "label"]
  - apiGroups: ["batch"]
    resources: ["jobs"]
    verbs: ["get", "list", "watch", "describe", "create", "edit", "delete", "annotate",
"patch", "label"]
  - apiGroups: ["extensions", "networking.k8s.io"]
    resources: ["ingresses"]
    verbs: ["get", "list", "watch", "describe", "create", "edit", "delete", "annotate",
"patch", "label"]
  - apiGroups: ["rbac.authorization.k8s.io"]
    resources: ["clusterroles", "clusterrolebindings", "roles", "rolebindings"]
    verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"deletecollection", "annotate", "patch", "label"]
  - apiGroups: [""]
    resources: ["persistentvolumeclaims"]
    verbs: ["get", "list", "watch", "describe", "create", "edit", "delete",
"deletecollection", "annotate", "patch", "label"]
  - apiGroups: ["kyverno.io"]
    resources: ["clusterpolicies"]
    verbs: ["create", "delete"]
---

```

```
kubectl apply -f emr-containers-cluster-role.yaml
```


Create Cluster Role Bindings

```
vim emr-containers-cluster-role-binding.yaml
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: emr-containers
subjects:
- kind: User
  name: emr-containers
  apiGroup: rbac.authorization.k8s.io
roleRef:
  kind: ClusterRole
  name: emr-containers
  apiGroup: rbac.authorization.k8s.io
---
```

```
kubectl apply -f emr-containers-cluster-role-binding.yaml
```

Provide Namespace access to the Amazon EMR on EKS service

Create two Kubernetes namespaces, one for User driver and executors, and another for System driver & executors, and enable Amazon EMR on EKS service access to submit Jobs in both User and System Namespaces. Follow the existing guide to provide access for each namespace, which is available at [Enable cluster access using aws-auth](#).

Step 3: Setup IAM Roles for user and system profile components

Third, you set up roles for specific components. A Lake Formation-enabled Spark Job has two components, User and System. The User driver and executors run in User namespace, and are tied to the JobExecutionRole that is passed in the StartJobRun API. The System driver and executors run in the System namespace, and are tied to the **QueryEngine** role.

Configure Query Engine role

The QueryEngine role is tied to the System Space Components, and would have permissions to assume the **JobExecutionRole** with **LakeFormationAuthorizedCaller** Session tag. The IAM Permissions Policy of Query Engine role is the following:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AssumeJobRoleWithSessionTagAccessForSystemDriver",
      "Effect": "Allow",
      "Action": [
        "sts:AssumeRole",
        "sts:TagSession"
      ],
      "Resource": [
        "arn:aws:iam::*:role/JobExecutionRole"
      ],
      "Condition": {
        "StringLike": {
          "aws:RequestTag/LakeFormationAuthorizedCaller": "EMR on EKS Engine"
        }
      }
    },
    {
      "Sid": "AssumeJobRoleWithSessionTagAccessForSystemExecutor",
      "Effect": "Allow",
      "Action": [
        "sts:AssumeRole"
      ],
      "Resource": [
        "arn:aws:iam::*:role/JobExecutionRole"
      ]
    },
    {
      "Sid": "CreateCertificateAccessForTLS",
      "Effect": "Allow",
      "Action": [
        "emr-containers:CreateCertificate"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

```
}
```

Configure the Trust policy of Query Engine role to trust the Kubernetes System namespace.

```
aws emr-containers update-role-trust-policy \  
  --cluster-name eks cluster \  
  --namespace eks system namespace \  
  --role-name query_engine_iam_role_name
```

For more information, see [Updating the role trust policy](#).

Configure the Job Execution Role

Lake Formation permissions control access to AWS Glue Data Catalog resources, Amazon S3 locations, and the underlying data at those locations. IAM permissions control access to the Lake Formation and AWS Glue APIs and resources. Although you might have the Lake Formation permission to access a table in the Data Catalog (SELECT), your operation fails if you don't have the IAM permission on the `glue:Get*` API operations.

IAM Permissions Policy of **JobExecutionRole**: The **JobExecution** Role should have the Policy Statements in its Permissions Policy.

JSON

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Sid": "GlueCatalogAccess",  
      "Effect": "Allow",  
      "Action": [  
        "glue:Get*",  
        "glue:Create*",  
        "glue:Update*"  
      ],  
      "Resource": [  
        "*"  
      ]  
    },  
  ],  
}
```

```

    "Sid": "LakeFormationAccess",
    "Effect": "Allow",
    "Action": [
        "lakeformation:GetDataAccess"
    ],
    "Resource": [
        "*"
    ]
  },
  {
    "Sid": "CreateCertificateAccessForTLS",
    "Effect": "Allow",
    "Action": [
        "emr-containers:CreateCertificate"
    ],
    "Resource": [
        "*"
    ]
  }
]
}

```

IAM Trust Policy for **JobExecutionRole**:

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "TrustQueryEngineRoleForSystemDriver",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:role/QueryEngineRole"
      },
      "Action": [
        "sts:AssumeRole",
        "sts:TagSession"
      ],
      "Condition": {
        "StringLike": {
          "aws:RequestTag/LakeFormationAuthorizedCaller": "EMR on EKS Engine"
        }
      }
    }
  ]
}

```

```

    }
  },
  {
    "Sid": "TrustQueryEngineRoleForSystemExecutor",
    "Effect": "Allow",
    "Principal": {
      "AWS": "arn:aws:iam::123456789012:role/QueryEngineRole"
    },
    "Action": "sts:AssumeRole"
  }
]
}

```

Configure the Trust Policy of Job execution Role to trust the Kubernetes user namespace:

```

aws emr-containers update-role-trust-policy \
  --cluster-name eks cluster \
  --namespace eks User namespace \
  --role-name job_execution_role_name

```

For more information, see [Update the trust policy of the job execution role](#).

Step 4: Setup security configuration

To run a Lake Formation-enabled job, you must create a security configuration.

```

aws emr-containers create-security-configuration \
  --name 'security-configuration-name' \
  --security-configuration '{
    "authorizationConfiguration": {
      "lakeFormationConfiguration": {
        "authorizedSessionTagValue": "SessionTag configured in LakeFormation",
        "secureNamespaceInfo": {
          "clusterId": "eks-cluster-name",
          "namespace": "system-namespace-name"
        },
        "queryEngineRoleArn": "query-engine-IAM-role-ARN"
      }
    }
  }'

```

Ensure that the Session Tag passed in the field **authorizedSessionTagValue** can authorize Lake Formation. Set the value to the one configured in Lake Formation, in [Step 1: Set up Lake Formation-based column, row, or cell-level permissions](#).

Step 5: Create a virtual cluster

Create a Amazon EMR on EKS virtual cluster with a security configuration.

```
aws emr-containers create-virtual-cluster \  
--name my-lf-enabled-vc \  
--container-provider '{  
  "id": "eks-cluster",  
  "type": "EKS",  
  "info": {  
    "eksInfo": {  
      "namespace": "user-namespace"  
    }  
  }  
' \  
--security-configuration-id SecurityConfigurationId
```

Ensure the **SecurityConfiguration** Id from the previous step is passed, so that the Lake Formation authorization configuration is applied to all Jobs running on the virtual cluster. For more information, see [Register the Amazon EKS cluster with Amazon EMR](#).

Step 6: Submit a Job in the FGAC Enabled VirtualCluster

The process for job submission is the same for both non Lake Formation and Lake Formation jobs. For more information, see [Submit a job run with StartJobRun](#).

The Spark Driver, Executor and Event Logs of the System Driver are stored in AWS Service Account's S3 Bucket for debugging. We recommend configuring a customer-managed KMS Key in the Job Run to encrypt all logs stored in the AWS service bucket. For more information about enabling log encryption, see [Encrypting Amazon EMR on EKS logs](#).

Considerations and limitations

Note the following considerations and limitations when you use Lake Formation with Amazon EMR on EKS:

- Amazon EMR on EKS supports fine-grained access control via Lake Formation only for Apache Hive, Apache Iceberg, Apache Hudi, and Delta table Formats. Apache Hive formats include Parquet, ORC, and xSV.
- `DynamicResourceAllocation` is enabled by default, and you can't turn off `DynamicResourceAllocation` for Lake Formation jobs. As `DRA spark.dynamicAllocation.maxExecutors` configuration's default value is infinity, please configure an appropriate value based on your workload.
- By default, `spark.dynamicAllocation.preallocateExecutors` is enabled in Amazon EMR Spark, which can cause excessive container churn when `spark.dynamicAllocation.initialExecutors` and `spark.dynamicAllocation.minExecutors` are not set. For recommended configurations to manage executor preallocation, see the [Performance](#) section in [Links to Amazon EMR on EKS best practices guides on GitHub](#).
- Lake Formation-enabled jobs don't support usage of customized EMR on EKS Images in System Driver and System Executors.
- You can only use Lake Formation with Spark jobs.
- EMR on EKS with Lake Formation only supports a single Spark session throughout a job.
- EMR on EKS with Lake Formation only supports cross-account table queries shared through resource links.
- The following aren't supported:
 - Resilient distributed datasets (RDD)
 - Spark streaming
 - Write with Lake Formation granted permissions
 - Access control for nested columns
- EMR on EKS blocks functionalities that might undermine the complete isolation of system driver, including the following:
 - UDTs, HiveUDFs, and any user-defined function that involves custom classes
 - Custom data sources
 - Supply of additional jars for Spark extension, connector, or metastore `ANALYZE TABLE` command
- To enforce access controls, `EXPLAIN PLAN` and DDL operations such as `DESCRIBE TABLE` don't expose restricted information.

- Amazon EMR on EKS restricts access to system driver Spark logs on Lake Formation-enabled jobs. Since the system driver runs with more access, events and logs that the system driver generates can include sensitive information. To prevent unauthorized users or code from accessing this sensitive data, EMR on EKS disabled access to system driver logs. For troubleshooting, contact AWS support.
- If you registered a table location with Lake Formation, the data access path goes through the Lake Formation stored credentials, regardless of the IAM permission for the EMR on EKS job execution role. If you misconfigure the role registered with the table location, jobs submitted that use the role with S3 IAM permission to the table location will fail.
- Writing to a Lake Formation table uses IAM permission rather than Lake Formation granted permissions. If your job execution role has the necessary S3 permissions, you can use it to run write operations.

The following are considerations and limitations when using Apache Iceberg:

- You can only use Apache Iceberg with session catalog and not arbitrarily named catalogs.
- Iceberg tables that are registered in Lake Formation only support the metadata tables history, metadata_log_entries, snapshots, files, manifests, and refs. Amazon EMR hides the columns that might have sensitive data, such as partitions, path, and summaries. This limitation doesn't apply to Iceberg tables that aren't registered in Lake Formation.
- Tables that you don't register in Lake Formation support all Iceberg stored procedures. The register_table and migrate procedures aren't supported for any tables.
- We recommend that you use Iceberg DataFrameWriterV2 instead of V1.

For more information, see [Understanding Amazon EMR on EKS concepts and terminology](#) and [Enable cluster access for Amazon EMR on EKS](#).

Disclaimer for data administrators

Note

When you grant access to Lake Formation resources to an IAM role for EMR on EKS, you must ensure the EMR cluster administrator or operator is a trusted administrator. This is particularly relevant for Lake Formation resources that are shared across multiple organizations and AWS accounts.

Responsibilities for EKS administrators

- The System namespace should be protected. No user or resource or entity or tooling would be allowed to have any Kubernetes RBAC permissions on the Kubernetes resources in the System namespace.
- No user or resource or entity except the EMR on EKS service should have access to CREATE access to POD, CONFIG_MAP and SECRET in the User namespace.
- System drivers and System executors contain sensitive data. So, Spark events, Spark driver logs, and Spark executor logs in the System namespace should not be forwarded to external log storage systems.

Troubleshooting

Logging

EMR on EKS uses Spark resources profiles to split job execution. Amazon EMR on EKS uses the user profile to run the code you supplied, while the system profile enforces Lake Formation policies. You can access the logs for the containers ran as the user profile by configuring the StartJobRun request with [MonitoringConfiguration](#).

Spark History Server

The Spark History Server have all Spark events generated from the user profile and redacted events generated from the system driver. You can see all of the containers from both the user and system drivers in the **Executors** tab. However, log links are available only for the user profile.

Job failed with insufficient Lake Formation permissions

Make sure that your job execution role has the permissions to run SELECT and DESCRIBE on the table that you are accessing.

Job with RDD execution failed

EMR on EKS currently doesn't support resilient distributed dataset (RDD) operations on Lake Formation-enabled jobs.

Unable to access data files in Amazon S3

Make sure you have registered the location of the data lake in Lake Formation.

Security validation exception

EMR on EKS detected a security validation error. Contact AWS support for assistance.

Sharing AWS Glue Data Catalog and tables across accounts

You can share databases and tables across accounts and still use Lake Formation. For more information, see [Cross-account data sharing in Lake Formation](#) and [How do I share AWS Glue Data Catalog and tables cross-account using AWS Lake Formation?](#)

Iceberg Job throwing initialization error not setting the AWS region

Message is the following:

```
25/02/25 13:33:19 ERROR SparkFGACEExceptionSanitizer: Client received error with id =  
b921f9e6-f655-491f-b8bd-b2842cdc20c7,  
reason = IllegalArgumentException, message = Cannot initialize  
LakeFormationAwsClientFactory, please set client.region to a valid aws region
```

Make sure the Spark configuration `spark.sql.catalog.catalog_name.client.region` is set to a valid region.

Iceberg Job throwing SparkUnsupportedOperationException

Message is the following:

```
25/02/25 13:53:15 ERROR SparkFGACEExceptionSanitizer: Client received error with id =  
921fef42-0800-448b-bef5-d283d1278ce0,  
reason = SparkUnsupportedOperationException, message = Either glue.id or glue.account-  
id is set with non-default account.  
Cross account access with fine-grained access control is only supported with AWS  
Resource Access Manager.
```

Make sure the Spark Configuration `spark.sql.catalog.catalog_name.glue.account-id` is set to a valid account id.

Iceberg Job fail with "403 Access Denied" during MERGE operation

Message is the following:

```
software.amazon.awssdk.services.s3.model.S3Exception: Access Denied (Service: S3,  
Status Code: 403,
```

```
...
at
software.amazon.awssdk.services.s3.DefaultS3Client.deleteObject(DefaultS3Client.java:3365)
at org.apache.iceberg.aws.s3.S3FileIO.deleteFile(S3FileIO.java:162)
at org.apache.iceberg.io.FileIO.deleteFile(FileIO.java:86)
at
org.apache.iceberg.io.RollingFileWriter.closeCurrentWriter(RollingFileWriter.java:129)
```

Disable S3 Delete operations in Spark by adding the following property. `--conf spark.sql.catalog.s3-table-name.s3.delete-enabled=false`.

Logging and monitoring

To detect incidents, receive alerts when incidents occur, and respond to them, use these options with Amazon EMR on EKS:

- Monitor Amazon EMR on EKS with AWS CloudTrail - [AWS CloudTrail](#) provides a record of actions taken by a user, role, or an AWS service in Amazon EMR on EKS. It captures calls from the Amazon EMR console and code calls to the Amazon EMR on EKS API operations as events. This allows you to determine the request that was made to Amazon EMR on EKS, the IP address from which the request was made, who made the request, when it was made, and additional details. For more information, see [Logging Amazon EMR on EKS API calls using AWS CloudTrail](#).
- Use CloudWatch Events with Amazon EMR on EKS - CloudWatch Events delivers a near real-time stream of system events that describe changes in AWS resources. CloudWatch Events becomes aware of operational changes as they occur, responds to them, and takes corrective action as necessary, by sending messages to respond to the environment, activating functions, making changes, and capturing state information. To use CloudWatch Events with Amazon EMR on EKS, create a rule that triggers on an Amazon EMR on EKS API call via CloudTrail. For more information, see [Monitor jobs with Amazon CloudWatch Events](#).

Encrypting Amazon EMR on EKS logs with managed storage

The sections that follow show you how to configure encryption for logs.

Enable encryption

To encrypt logs in managed storage with your own KMS key, use the following configuration when you submit a job run.

```
"monitoringConfiguration": {
    "managedLogs": {
        "allowAWSToRetainLogs": "ENABLED",
        "encryptionKeyArn": "KMS key arn"
    },
    "persistentAppUI": "ENABLED"
}
```

The `allowAWSToRetainLogs` configuration allows AWS to retain system namespace logs when running a job using Native FGAC. The `persistentAppUI` configuration allows AWS to save event logs which are used to generate the Spark UI. The `encryptionKeyArn` is used to specify the KMS key ARN you want to use to encrypt the logs stored by AWS.

Required permissions for log encryption

The user who submits the job or views the Spark UI must be allowed the actions `kms:DescribeKey`, `kms:GenerateDataKey`, and `kms:Decrypt` for the encryption key. These permissions are used to verify the validity of the key and check that the user has the necessary permissions to read and write logs encrypted with the KMS key. If the user who submits the job lacks the necessary key permissions, Amazon EMR on EKS rejects the job run submission.

Example IAM Policy for Role Used to Call `StartJobRun`

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "emr-containers:StartJobRun"
      ],
      "Resource": [
        "*"
      ],
      "Effect": "Allow",
      "Sid": "AllowEMRCONTAINERSStartjobrun"
    },
    {
      "Action": [
        "kms:DescribeKey",
```

```

        "kms:Decrypt",
        "kms:GenerateDataKey"
    ],
    "Resource": [
        "arn:aws:kms:*:*:key/key-id"
    ],
    "Effect": "Allow",
    "Sid": "AllowKMSDescribekey"
}
]
}

```

You must also configure the KMS key to allow the `elasticmapreduce.amazonaws.com` and `elasticmapreduce.amazonaws.com` Service Principals to `kms:GenerateDataKey` and `kms:Decrypt`. This allows EMR to read and write logs encrypted with the KMS key to managed storage.

Example KMS Key Policy

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:DescribeKey"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringLike": {
          "kms:viaService": "emr-containers.*.amazonaws.com"
        }
      },
      "Sid": "AllowKMSDescribekey"
    },
    {

```

```

    "Effect": "Allow",
    "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKey"
    ],
    "Resource": [
        "*"
    ],
    "Condition": {
        "StringLike": {
            "kms:viaService": "emr-containers.*.amazonaws.com",
            "kms:EncryptionContext:aws:emr-containers:virtualClusterId": "virtual cluster id"
        }
    },
    "Sid": "AllowKMSDecryptGenerate"
},
{
    "Effect": "Allow",
    "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKey"
    ],
    "Resource": [
        "*"
    ],
    "Condition": {
        "StringLike": {
            "kms:EncryptionContext:aws:emr-containers:virtualClusterId": "virtual cluster id"
        },
        "ArnLike": {
            "aws:SourceArn": "arn:aws:emr-containers:*:*:/virtualclusters/virtual_cluster_id"
        }
    },
    "Sid": "AllowKMSDecryptService"
}
]
}

```

As a security best practice, we recommend that you add the `kms:viaService`, `kms:EncryptionContext`, and `aws:SourceArn` conditions. These conditions help ensure the key is only used by Amazon EMR on EKS and only used for logs generated from jobs running in a specific virtual cluster.

Logging Amazon EMR on EKS API calls using AWS CloudTrail

Amazon EMR on EKS is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Amazon EMR on EKS. CloudTrail captures all API calls for Amazon EMR on EKS as events. The calls captured include calls from the Amazon EMR on EKS console and code calls to the Amazon EMR on EKS API operations. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon EMR on EKS. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Amazon EMR on EKS, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, see the [AWS CloudTrail User Guide](#).

Amazon EMR on EKS information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When activity occurs in Amazon EMR on EKS, that activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see [Viewing events with CloudTrail Event history](#).

For an ongoing record of events in your AWS account, including events for Amazon EMR on EKS, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- [Overview for creating a trail](#)
- [CloudTrail supported services and integrations](#)
- [Configuring Amazon SNS notifications for CloudTrail](#)
- [Receiving CloudTrail log files from multiple regions](#) and [Receiving CloudTrail log files from multiple accounts](#)

All Amazon EMR on EKS actions are logged by CloudTrail and are documented in [Amazon EMR on EKS API documentation](#). For example, calls to the `CreateVirtualCluster`, `StartJobRun` and `ListJobRuns` actions generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the [CloudTrail user Identity element](#).

Understanding Amazon EMR on EKS log file entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the [ListJobRuns](#) action.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE:admin",
    "arn": "arn:aws:sts::012345678910:assumed-role/Admin/admin",
    "accountId": "012345678910",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",
        "arn": "arn:aws:iam::012345678910:role/Admin",
        "accountId": "012345678910",
        "userName": "Admin"
      }
    }
  },
```



```

    "webIdFederationData": {},
    "attributes": {
      "mfaAuthenticated": "false",
      "creationDate": "2020-11-04T21:49:36Z"
    }
  },
  "eventTime": "2020-11-04T21:52:58Z",
  "eventSource": "emr-containers.amazonaws.com",
  "eventName": "ListJobRuns",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "203.0.113.1",
  "userAgent": "aws-cli/1.11.167 Python/2.7.10 Darwin/16.7.0 botocore/1.7.25",
  "requestParameters": {
    "virtualClusterId": "1K48XXXXXXHCB"
  },
  "responseElements": null,
  "requestID": "890b8639-e51f-11e7-b038-EXAMPLE",
  "eventID": "874f89fa-70fc-4798-bc00-EXAMPLE",
  "readOnly": true,
  "eventType": "AwsApiCall",
  "recipientAccountId": "012345678910"
}

```

Using Amazon S3 Access Grants with Amazon EMR on EKS

S3 Access Grants overview for Amazon EMR on EKS

With Amazon EMR releases 6.15.0 and higher, Amazon S3 Access Grants provide a scalable access control solution that you can use to augment access to your Amazon S3 data from Amazon EMR on EKS. If you have a complex or large permission configuration for your S3 data, you can use Access Grants to scale S3 data permissions for users, roles, and applications.

Use S3 Access Grants to augment access to Amazon S3 data beyond the permissions granted by the runtime role or the IAM roles that are attached to the identities with access to your Amazon EMR on EKS cluster.

For more information, see [Managing access with S3 Access Grants for Amazon EMR](#) in the *Amazon EMR Management Guide* and [Managing access with S3 Access Grants](#) in the *Amazon Simple Storage Service User Guide*.

This page describes the requirements to run a Spark job in Amazon EMR on EKS with S3 Access Grants integration. With Amazon EMR on EKS, S3 Access Grants requires an additional IAM policy statement in the execution role for your job, and an additional override configuration for the `StartJobRun` API. For steps to set up S3 Access Grants with other Amazon EMR deployments, see the following documentation:

- [Using S3 Access Grants with Amazon EMR](#)
- [Using S3 Access Grants with EMR Serverless](#)

Launch an Amazon EMR on EKS cluster with S3 Access Grants for data management

You can enable S3 Access Grants on Amazon EMR on EKS and launch a Spark job. When your application makes a request for S3 data, Amazon S3 provides temporary credentials that are scoped to the specific bucket, prefix, or object.

1. Set up a job execution role for your Amazon EMR on EKS cluster. Include the required IAM permissions that you need to run Spark jobs, `s3:GetDataAccess` and `s3:GetAccessGrantsInstanceForPrefix`:

```
{
  "Effect": "Allow",
  "Action": [
    "s3:GetDataAccess",
    "s3:GetAccessGrantsInstanceForPrefix"
  ],
  "Resource": [
    //LIST ALL INSTANCE ARNS THAT THE ROLE IS ALLOWED TO QUERY
    "arn:aws_partition:s3:Region:account-id1:access-grants/default",
    "arn:aws_partition:s3:Region:account-id2:access-grants/default"
  ]
}
```

Note

If you specify IAM roles that for job execution that have any additional permissions to access S3 directly, then users might be able to access data regardless of the permissions that you define in S3 Access Grants

2. Submit a job to your Amazon EMR on EKS cluster with an Amazon EMR release label of 6.15 or higher and the `emrfs-site` classification, as the following example shows. Replace the values in *red text* with the appropriate values for your usage scenario.

```
{
  "name": "myjob",
  "virtualClusterId": "123456",
  "executionRoleArn": "iam_role_name_for_job_execution",
  "releaseLabel": "emr-7.13.0-latest",
  "jobDriver": {
    "sparkSubmitJobDriver": {
      "entryPoint": "entryPoint_location",
      "entryPointArguments": ["argument1", "argument2"],
      "sparkSubmitParameters": "--class main_class"
    }
  },
  "configurationOverrides": {
    "applicationConfiguration": [
      {
        "classification": "emrfs-site",
        "properties": {
          "fs.s3.s3AccessGrants.enabled": "true",
          "fs.s3.s3AccessGrants.fallbackToIAM": "false"
        }
      }
    ]
  }
}
```

S3 Access Grants considerations with Amazon EMR on EKS

For important support, compatibility, and behavioral information when you use Amazon S3 Access Grants with Amazon EMR on EKS, see [S3 Access Grants considerations with Amazon EMR](#) in the *Amazon EMR Management Guide*.

Compliance validation for Amazon EMR on EKS

Third-party auditors assess the security and compliance of Amazon EMR on EKS as part of multiple AWS compliance programs. These include SOC, PCI, FedRAMP, HIPAA, and others.

Resilience in Amazon EMR on EKS

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](#).

In addition to the AWS global infrastructure, Amazon EMR on EKS offers integration with Amazon S3 through EMRFS to help support your data resiliency and backup needs.

Infrastructure security in Amazon EMR on EKS

As a managed service, Amazon EMR 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 Amazon EMR 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.

Configuration and vulnerability analysis

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:

- [Compliance validation for Amazon EMR on EKS](#)
- [Shared Responsibility Model](#)

- [Amazon Web Services: Overview of Security Processes](#) (whitepaper)

Connect to Amazon EMR on EKS Using an interface VPC endpoint

You can connect directly to Amazon EMR on EKS using [Interface VPC endpoints \(AWS PrivateLink\)](#) in your Virtual Private Cloud (VPC) instead of connecting over the internet. When you use an interface VPC endpoint, communication between your VPC and Amazon EMR on EKS is conducted entirely within the AWS network. Each VPC endpoint is represented by one or more [Elastic network interfaces \(ENIs\)](#) with private IP addresses in your VPC subnets.

The interface VPC endpoint connects your VPC directly to Amazon EMR on EKS without an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. The instances in your VPC don't need public IP addresses to communicate with the Amazon EMR on EKS API.

You can create an interface VPC endpoint to connect to Amazon EMR on EKS using the AWS Management Console or AWS Command Line Interface (AWS CLI) commands. For more information, see [Creating an Interface Endpoint](#).

After you create an interface VPC endpoint, if you enable private DNS hostnames for the endpoint, the default Amazon EMR on EKS endpoint resolves to your VPC endpoint. The default service name endpoint for Amazon EMR on EKS is in the following format.

```
emr-containers.Region.amazonaws.com
```

If you do not enable private DNS hostnames, Amazon VPC provides a DNS endpoint name that you can use in the following format.

```
VPC_Endpoint_ID.emr-containers.Region.vpce.amazonaws.com
```

For more information, see [Interface VPC Endpoints \(AWS PrivateLink\)](#) in the Amazon VPC User Guide. Amazon EMR on EKS supports making calls to all of its [API Actions](#) inside your VPC.

You can attach VPC endpoint policies to a VPC endpoint to control access for IAM principals. You can also associate security groups with a VPC endpoint to control inbound and outbound access based on the origin and destination of network traffic, such as a range of IP addresses. For more information, see [Controlling Access to Services with VPC Endpoints](#).

Create a VPC Endpoint Policy for Amazon EMR on EKS

You can create a policy for Amazon VPC endpoints for Amazon EMR on EKS to specify the following:

- The principal that can or cannot perform actions
- The actions that can be performed
- The resources on which actions can be performed

For more information, see [Controlling Access to Services with VPC Endpoints](#) in the Amazon VPC User Guide.

Example VPC Endpoint Policy to Deny All Access From a Specified AWS Account

The following VPC endpoint policy denies AWS account **123456789012** all access to resources using the endpoint.

```
{
  "Statement": [
    {
      "Action": "*",
      "Effect": "Allow",
      "Resource": "*",
      "Principal": "*"
    },
    {
      "Action": "*",
      "Effect": "Deny",
      "Resource": "*",
      "Principal": {
        "AWS": [
          "123456789012"
        ]
      }
    }
  ]
}
```

Example VPC Endpoint Policy to Allow VPC Access Only to a Specified IAM Principal (User)

The following VPC endpoint policy allows full access only to the IAM user *lijuan* in AWS account *123456789012*. All other IAM principals are denied access using the endpoint.

```
{
  "Statement": [
    {
      "Action": "*",
      "Effect": "Allow",
      "Resource": "*",
      "Principal": {
        "AWS": [
          "arn:aws:iam::123456789012:user/lijuan"
        ]
      }
    }
  ]
}
```

Example VPC Endpoint Policy to Allow Read-Only Amazon EMR on EKS Operations

The following VPC endpoint policy allows only AWS account *123456789012* to perform the specified Amazon EMR on EKS actions.

The actions specified provide the equivalent of read-only access for Amazon EMR on EKS. All other actions on the VPC are denied for the specified account. All other accounts are denied any access. For a list of Amazon EMR on EKS actions, see [Actions, Resources, and Condition Keys for Amazon EMR on EKS](#).

```
{
  "Statement": [
    {
      "Action": [
        "emr-containers:DescribeJobRun",
        "emr-containers:DescribeVirtualCluster",
        "emr-containers:ListJobRuns",
        "emr-containers:ListTagsForResource",
        "emr-containers:ListVirtualClusters"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

```

        "Principal": {
            "AWS": [
                "123456789012"
            ]
        }
    ]
}

```

Example VPC Endpoint Policy Denying Access to a Specified Virtual Cluster

The following VPC endpoint policy allows full access for all accounts and principals, but denies any access for AWS account **123456789012** to actions performed on the virtual cluster with cluster ID **A1B2CD34EF5G**. Other Amazon EMR on EKS actions that don't support resource-level permissions for virtual clusters are still allowed. For a list of Amazon EMR on EKS actions and their corresponding resource type, see [Actions, Resources, and Condition Keys for Amazon EMR on EKS](#) in the *AWS Identity and Access Management User Guide*.

```

{
    "Statement": [
        {
            "Action": "*",
            "Effect": "Allow",
            "Resource": "*",
            "Principal": "*"
        },
        {
            "Action": "*",
            "Effect": "Deny",
            "Resource": "arn:aws:emr-containers:us-west-2:123456789012:/virtualclusters/A1B2CD34EF5G",
            "Principal": {
                "AWS": [
                    "123456789012"
                ]
            }
        }
    ]
}

```


Set up cross-account access for Amazon EMR on EKS

You can set up cross-account access for Amazon EMR on EKS. Cross-account access enables users from one AWS account to run Amazon EMR on EKS jobs and access the underlying data that belongs to another AWS account.

Prerequisites

To set up cross-account access for Amazon EMR on EKS, you'll complete tasks while signed in to the following AWS accounts:

- **AccountA** - An AWS account where you have created an Amazon EMR on EKS virtual cluster by registering Amazon EMR with a namespace on an EKS cluster.
- **AccountB** - An AWS account that contains an Amazon S3 bucket or a DynamoDB table that you want your Amazon EMR on EKS jobs to access.

You must have the following ready in your AWS accounts before setting up cross-account access:

- An Amazon EMR on EKS virtual cluster in AccountA where you want to run jobs.
- A job execution role in AccountA that has the required permissions to run jobs in the virtual cluster. For more information, see [Create a job execution role](#) and [Using job execution roles with Amazon EMR on EKS](#).

How to access a cross-account Amazon S3 bucket or DynamoDB table

To set up cross-account access for Amazon EMR on EKS, complete the following steps.

1. Create an Amazon S3 bucket, `cross-account-bucket`, in AccountB. For more information, see [Creating a bucket](#). If you want to have cross-account access to DynamoDB, you can also create a DynamoDB table in AccountB. For more information, see [Creating a DynamoDB table](#).
2. Create a `Cross-Account-Role-B` IAM role in AccountB that can access the `cross-account-bucket`.
 1. Sign in to the IAM console.
 2. Choose **Roles** and create a new role: `Cross-Account-Role-B`. For more information about how to create IAM roles, see [Creating IAM roles](#) in the IAM user Guide.

3. Create an IAM policy that specifies the permissions for Cross-Account-Role-B to access the cross-account-bucket S3 bucket, as the following policy statement demonstrates. Then attach the IAM policy to Cross-Account-Role-B. For more information, see [Creating a New Policy](#) in the IAM user Guide.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:*"
      ],
      "Resource": [
        "arn:aws:s3:::cross-account-bucket",
        "arn:aws:s3:::cross-account-bucket/*"
      ],
      "Sid": "AllowS3"
    }
  ]
}
```

If DynamoDB access is required, create an IAM policy that specifies permissions to access the cross-account DynamoDB table. Then attach the IAM policy to Cross-Account-Role-B. For more information, see [Create a DynamoDB table](#) in the IAM user guide.

Following is a policy to access a DynamoDB table, CrossAccountTable.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "dynamodb:*"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:*:table/CrossAccountTable"
      ]
    }
  ]
}
```

```

    ],
    "Sid": "AllowDYNAMODB"
  }
]
}

```

3. Edit the trust relationship for the Cross-Account-Role-B role.

1. To configure the trust relationship for the role, choose the **Trust Relationships** tab in the IAM console for the role created in Step 2: Cross-Account-Role-B.
2. Select **Edit Trust Relationship**.
3. Add the following policy document, which allows Job-Execution-Role-A in AccountA to assume this Cross-Account-Role-B role.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowSTSAssumerole",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:role/Job-Execution-Role-A"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}

```

4. Grant Job-Execution-Role-A in AccountA with - STS Assume role permission to assume Cross-Account-Role-B.

1. In the IAM console for AWS account AccountA, select Job-Execution-Role-A.
2. Add the following policy statement to the Job-Execution-Role-A to allow the AssumeRole action on the Cross-Account-Role-B role.

JSON

```

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

```

```
{
  "Effect": "Allow",
  "Action": [
    "sts:AssumeRole"
  ],
  "Resource": [
    "arn:aws:iam::*:role/Cross-Account-Role-B"
  ],
  "Sid": "AllowSTSAssumerole"
}
]
```

5. For Amazon S3 access, set the following spark-submit parameters (spark conf) while submitting the job to Amazon EMR on EKS.

Note

By default, EMRFS uses the job execution role to access the S3 bucket from the job. But when customAWSCredentialsProvider is set to AssumeRoleAWSCredentialsProvider, EMRFS uses the corresponding role that you specify with ASSUME_ROLE_CREDENTIALS_ROLE_ARN instead of the Job-Execution-Role-A for Amazon S3 access.

- --conf
spark.hadoop.fs.s3.customAWSCredentialsProvider=com.amazonaws.emr.AssumeRoleAWSCredentialsProvider
- --conf
spark.kubernetes.driverEnv.ASSUME_ROLE_CREDENTIALS_ROLE_ARN=arn:aws:iam::**AccountA**:role/Cross-Account-Role-B \
- --conf
spark.executorEnv.ASSUME_ROLE_CREDENTIALS_ROLE_ARN=arn:aws:iam::**AccountB**:role/Cross-Account-Role-B \

Note

You must set ASSUME_ROLE_CREDENTIALS_ROLE_ARN for both executor and driver env in the job spark configuration.

For DynamoDB cross-account access, you must set `--conf`

`spark.dynamodb.customAWSCredentialsProvider=com.amazonaws.emr.AssumeRoleAWSCr`

6. Run the Amazon EMR on EKS job with cross-account access, as the following example demonstrates.

```
aws emr-containers start-job-run \
--virtual-cluster-id 123456 \
--name myjob \
--execution-role-arn execution-role-arn \
--release-label emr-6.2.0-latest \
--job-driver '{"sparkSubmitJobDriver": {"entryPoint": "entryPoint_location",
"entryPointArguments": ["arguments_list"], "sparkSubmitParameters": "--class
<main_class> --conf spark.executor.instances=2 --conf spark.executor.memory=2G
--conf spark.executor.cores=2 --conf spark.driver.cores=1 --conf
spark.hadoop.fs.s3.customAWSCredentialsProvider=com.amazonaws.emr.AssumeRoleAWSCredentials
--conf
spark.kubernetes.driverEnv.ASSUME_ROLE_CREDENTIALS_ROLE_ARN=arn:aws:iam::AccountB:role/
Cross-Account-Role-B --conf
spark.executorEnv.ASSUME_ROLE_CREDENTIALS_ROLE_ARN=arn:aws:iam::AccountB:role/
Cross-Account-Role-B"}} ' \
--configuration-overrides '{"applicationConfiguration": [{"classification":
"spark-defaults", "properties": {"spark.driver.memory": "2G"}}],
"monitoringConfiguration": {"cloudWatchMonitoringConfiguration":
{"logGroupName": "log_group_name", "logStreamNamePrefix": "log_stream_prefix"},
"persistentAppUI": "ENABLED", "s3MonitoringConfiguration": {"logUri": "s3://
my_s3_log_location" }}}
```

Tagging your Amazon EMR on EKS resources

To help you manage your Amazon EMR on EKS resources, you can assign your own metadata to each resource using tags. This topic provides an overview of the tags function and shows you how to create tags.

Topics

- [Tag basics](#)
- [Tag your resources](#)
- [Tag restrictions](#)
- [Work with tags using the AWS CLI and the Amazon EMR on EKS API](#)

Tag basics

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

Tags enable you to categorize your AWS resources by attributes such as purpose, owner, or environment. When you have many resources of the same type, you can quickly identify a specific resource based on the tags you've assigned to it. For example, you can define a set of tags for your Amazon EMR on EKS clusters to help you track each cluster's owner and stack level. We recommend that you devise a consistent set of tag keys for each resource type. You can then search and filter the resources based on the tags that you add.

Tags are not automatically assigned to your resources. After you add a tag, you can edit tag keys and values or remove tags from a resource at any time. If you delete a resource, any tags for the resource are also deleted.

Tags don't have any semantic meaning to Amazon EMR on EKS and are interpreted strictly as a string of characters.

A tag value can be an empty string, but not null. A tag key cannot be an empty string. If you add a tag that has the same key as an existing tag on that resource, the new value overwrites the earlier value.

If you use AWS Identity and Access Management (IAM), you can control which users in your AWS account have permission to manage tags.

For tag-based access control policy examples, see [Policies for tag-based access control](#).

Tag your resources

You can tag new or existing virtual clusters and job runs that are in active states. The active states for job runs include: PENDING, SUBMITTED, RUNNING, and CANCEL_PENDING. The active states for virtual clusters include: RUNNING, TERMINATING and ARRESTED. For more information, see [Job run states](#) and [Virtual cluster states](#).

When a virtual cluster is terminated, tags are cleaned and no longer accessible.

If you're using the Amazon EMR on EKS API, the AWS CLI, or an AWS SDK, you can apply tags to new resources using the tags parameter on the relevant API action. You can apply tags to existing resources using the TagResource API action.

You can use some resource-creating actions to specify tags for a resource when the resource is created. In this case, if tags cannot be applied while the resource is being created, the resource fails to be created. This mechanism ensures that resources you intended to tag on creation are either created with specified tags or not created at all. If you tag resources at the time of creation, you don't need to run custom tagging scripts after creating a resource.

The following table describes the Amazon EMR on EKS resources that can be tagged.

Resource	Supports tags	Supports tag propagation	Supports tagging on creation (Amazon EMR on EKS API, AWS CLI, and AWS SDK)	API for creation (tags can be added during creation)
Virtual cluster	Yes	No. Tags associated with a virtual cluster do not propagate to job runs submitted to that virtual cluster.	Yes	CreateVirtualCluster

Resource	Supports tags	Supports tag propagation	Supports tagging on creation (Amazon EMR on EKS API, AWS CLI, and AWS SDK)	API for creation (tags can be added during creation)
Job runs	Yes	No	Yes	StartJobRun

Tag restrictions

The following basic restrictions apply to tags:

- Maximum number of tags per resource – 50
- For each resource, each tag key must be unique, and each tag key can have only one value.
- Maximum key length – 128 Unicode characters in UTF-8
- Maximum value length – 256 Unicode characters in UTF-8
- If your tagging schema is used across multiple AWS services and resources, remember that other services may have restrictions on allowed characters. Generally allowed characters are letters, numbers, spaces representable in UTF-8, and the following characters: + - = . _ : / @.
- Tag keys and values are case sensitive.
- A tag value can be an empty string, but not null. A tag key cannot be an empty string.
- Don't use `aws :`, `AWS :`, or any upper or lowercase combination of such as a prefix for either keys or values. These are reserved only for AWS use.

Work with tags using the AWS CLI and the Amazon EMR on EKS API

Use the following AWS CLI commands or Amazon EMR on EKS API operations to add, update, list, and delete the tags for your resources.

Task	AWS CLI	API action
Add or overwrite one or more tags	tag-resource	TagResource
List tags for a resource	list-tags-for-resource	ListTagsForResource
Delete one or more tags	untag-resource	UntagResource

The following examples show how to tag or untag resources using the AWS CLI.

Example 1: Tag an existing virtual cluster

The following command tags an existing virtual cluster.

```
aws emr-containers tag-resource --resource-arn resource_ARN --tags team=devs
```

Example 2: Untag an existing virtual cluster

The following command deletes a tag from an existing virtual cluster.

```
aws emr-containers untag-resource --resource-arn resource_ARN --tag-keys tag_key
```

Example 3: List tags for a resource

The following command lists the tags associated with an existing resource.

```
aws emr-containers list-tags-for-resource --resource-arn resource_ARN
```

Troubleshooting for Amazon EMR on EKS

This section describes how to troubleshoot problems with Amazon EMR on EKS. For information about how to troubleshoot general problems with Amazon EMR, see [Troubleshoot a cluster](#) in the *Amazon EMR Management Guide*.

Topics

- [Troubleshooting jobs that use PersistentVolumeClaims \(PVC\)](#)
- [Troubleshooting Amazon EMR on EKS vertical autoscaling](#)
- [Troubleshooting Amazon EMR on EKS Spark operator](#)

Troubleshooting jobs that use PersistentVolumeClaims (PVC)

If you need to create, list, or delete PersistentVolumeClaims (PVC) for a job but don't add PVC permissions to the default Kubernetes role *emr-containers*, the job fails when you submit it. Without these permissions, the *emr-containers* role can't create necessary roles for the Spark driver or Spark client. It isn't enough to add permissions to the Spark driver or client roles, as suggested by error messages. The *emr-containers* primary role must include the required permissions also. This section explains how to add the required permissions to the *emr-containers* primary role.

Verification

To verify whether or not your *emr-containers* role has the necessary permissions, set the `NAMESPACE` variable with your own value and then run the following command:

```
export NAMESPACE=YOUR_VALUE
kubectl describe role emr-containers -n ${NAMESPACE}
```

In addition, to verify whether the Spark and client roles have the necessary permissions, run the following command:

```
kubectl describe role emr-containers-role-spark-driver -n ${NAMESPACE}
kubectl describe role emr-containers-role-spark-client -n ${NAMESPACE}
```

If the permissions aren't there, proceed with the patch, as follows.

Patch

1. If the jobs without the permissions are currently running, stop these jobs.
2. Create a file named *RBAC_Patch.py* as follows:

```
import os
import subprocess as sp
import tempfile as temp
import json
import argparse
import uuid

def delete_if_exists(dictionary: dict, key: str):
    if dictionary.get(key, None) is not None:
        del dictionary[key]

def doTerminalCmd(cmd):
    with temp.TemporaryFile() as f:
        process = sp.Popen(cmd, stdout=f, stderr=f)
        process.wait()
        f.seek(0)
        msg = f.read().decode()
    return msg

def patchRole(roleName, namespace, extraRules, skipConfirmation=False):
    cmd = f"kubectl get role {roleName} -n {namespace} --output json".split(" ")
    msg = doTerminalCmd(cmd)
    if "(NotFound)" in msg and "Error" in msg:
        print(msg)
        return False
    role = json.loads(msg)
    rules = role["rules"]
    rulesToAssign = extraRules[::]
    passedRules = []
    for rule in rules:
        apiGroups = set(rule["apiGroups"])
        resources = set(rule["resources"])
        verbs = set(rule["verbs"])
        for extraRule in extraRules:
            passes = 0
            apiGroupsExtra = set(extraRule["apiGroups"])
            resourcesExtra = set(extraRule["resources"])
            verbsExtra = set(extraRule["verbs"])
```

```

        passes += len(apiGroupsExtra.intersection(apiGroups)) >=
len(apiGroupsExtra)
        passes += len(resourcesExtra.intersection(resources)) >=
len(resourcesExtra)
        passes += len(verbsExtra.intersection(verbs)) >= len(verbsExtra)
        if passes >= 3:
            if extraRule not in passedRules:
                passedRules.append(extraRule)
                if extraRule in rulesToAssign:
                    rulesToAssign.remove(extraRule)
            break
    prompt_text = "Apply Changes?"
    if len(rulesToAssign) == 0:
        print(f"The role {roleName} seems to already have the necessary
permissions!")
        prompt_text = "Proceed anyways?"
    for ruleToAssign in rulesToAssign:
        role["rules"].append(ruleToAssign)
    delete_if_exists(role, "creationTimestamp")
    delete_if_exists(role, "resourceVersion")
    delete_if_exists(role, "uid")
    new_role = json.dumps(role, indent=3)
    uid = uuid.uuid4()
    filename = f"Role-{roleName}-New_Permissions-{uid}-TemporaryFile.json"
    try:
        with open(filename, "w+") as f:
            f.write(new_role)
            f.flush()
        prompt = "y"
        if not skipConfirmation:
            prompt = input(
                doTerminalCmd(f"kubectl diff -f {filename}".split(" ")) +
f"\n{prompt_text} y/n: "
            ).lower().strip()
            while prompt != "y" and prompt != "n":
                prompt = input("Please make a valid selection. y/n:
").lower().strip()
            if prompt == "y":
                print(doTerminalCmd(f"kubectl apply -f {filename}".split(" ")))
    except Exception as e:
        print(e)
    os.remove(f"./{filename}")

if __name__ == '__main__':

```

```

parser = argparse.ArgumentParser()
parser.add_argument("-n", "--namespace",
                    help="Namespace of the Role. By default its the
VirtualCluster's namespace",
                    required=True,
                    dest="namespace"
                    )

parser.add_argument("-p", "--no-prompt",
                    help="Applies the patches without asking first",
                    dest="no_prompt",
                    default=False,
                    action="store_true"
                    )
args = parser.parse_args()

emrRoleRules = [
    {
        "apiGroups": [""],
        "resources": ["persistentvolumeclaims"],
        "verbs": ["list", "create", "delete", "patch"]
    }
]

driverRoleRules = [
    {
        "apiGroups": [""],
        "resources": ["persistentvolumeclaims"],
        "verbs": ["list", "create", "delete", "patch", "deletecollection"]
    },
    {
        "apiGroups": [""],
        "resources": ["services"],
        "verbs": ["get", "list", "describe", "create", "delete", "watch",
"deletecollection"]
    },
    {
        "apiGroups": [""],
        "resources": ["configmaps", "pods"],
        "verbs": ["deletecollection"]
    }
]

```

```
clientRoleRules = [
    {
        "apiGroups": [""],
        "resources": ["persistentvolumeclaims"],
        "verbs": ["list", "create", "delete", "patch"]
    }
]

patchRole("emr-containers", args.namespace, emrRoleRules, args.no_prompt)
patchRole("emr-containers-role-spark-driver", args.namespace, driverRoleRules,
args.no_prompt)
patchRole("emr-containers-role-spark-client", args.namespace, clientRoleRules,
args.no_prompt)
```

3. Run the Python script:

```
python3 RBAC_Patch.py -n ${NAMESPACE}
```

4. A `kubectldiff` between the new permissions and the old ones appears. Press `y` to patch the role.

5. Verify the three roles with additional permissions as follows:

```
kubectl describe role -n ${NAMESPACE}
```

6. Run the python script:

```
python3 RBAC_Patch.py -n ${NAMESPACE}
```

7. After running the command, it will show a `kubectldiff` between the new permissions and the old ones. Press `y` to patch the role.

8. Verify the three roles with additional permissions:

```
kubectl describe role -n ${NAMESPACE}
```

9. Submit the job again.

Manual patch

If the permission that your application requires applies to something other than the PVC rules, you can manually add Kubernetes permissions for your Amazon EMR virtual cluster as needed.

Note

The role *emr-containers* is a primary role. This means that it must provide all the necessary permissions before you can change your underlying driver or client roles.

1. Download the current permissions into yaml files by running the commands below:

```
kubectl get role -n ${NAMESPACE} emr-containers -o yaml >> emr-containers-role-patch.yaml
kubectl get role -n ${NAMESPACE} emr-containers-role-spark-driver -o yaml >> driver-role-patch.yaml
kubectl get role -n ${NAMESPACE} emr-containers-role-spark-client -o yaml >> client-role-patch.yaml
```

2. Based on the permission your application requires, edit each file and add additional rules such as the following:

- `emr-containers-role-patch.yaml`

```
- apiGroups:
  - ""
  resources:
  - persistentvolumeclaims
  verbs:
  - list
  - create
  - delete
  - patch
```

- `driver-role-patch.yaml`

```
- apiGroups:
  - ""
  resources:
  - persistentvolumeclaims
  verbs:
  - list
  - create
  - delete
  - patch
  - deletecollection
```

```
- apiGroups:
  - ""
  resources:
  - services
  verbs:
  - get
  - list
  - describe
  - create
  - delete
  - watch
  - deletecollection
- apiGroups:
  - ""
  resources:
  - configmaps
  - pods
  verbs:
  - deletecollection
```

- `client-role-patch.yaml`

```
- apiGroups:
  - ""
  resources:
  - persistentvolumeclaims
  verbs:
  - list
  - create
  - delete
  - patch
```

3. Remove the following attributes with their values. This is necessary to apply the update.

- `creationTimestamp`
- `resourceVersion`
- `uid`

4. Finally, run the patch:

```
kubectl apply -f emr-containers-role-patch.yaml
kubectl apply -f driver-role-patch.yaml
kubectl apply -f client-role-patch.yaml
```


Troubleshooting Amazon EMR on EKS vertical autoscaling

Refer to the following sections if you encounter problems when you set up the Amazon EMR on EKS vertical autoscaling operator on an Amazon EKS cluster with Operator Lifecycle Manager. For more information including steps to complete the installation, see [Using vertical autoscaling with Amazon EMR Spark jobs](#).

403 Forbidden error

If you followed the steps in [Install the Operator Lifecycle Manager \(OLM\) on your Amazon EKS cluster](#), ran the `olm status` command, and it returned a 403 Forbidden error like the one below, you might not have obtained the authentication tokens to the Amazon ECR repository for the operator.

To resolve this issue, repeat the step in [Install the Amazon EMR on EKS vertical autoscaling operator](#) to obtain the tokens. Then, try the installation again.

```
Error: FATA[0002] Failed to run bundle: pull bundle image: error pulling image IMAGE.  
error resolving name : unexpected status code [manifests latest]: 403 Forbidden
```

Kubernetes namespace not found

When you [set up the Amazon EMR on EKS vertical autoscaling operator](#) on an Amazon EKS cluster, you might get a namespaces not found error like the one shown here:

```
FATA[0020] Failed to run bundle: create catalog: error creating catalog source:  
namespaces "NAME" not found.
```

If the namespace that you specify doesn't exist, OLM won't install the vertical autoscaling operator. To resolve this issue, use the following command to create the namespace. Then, try the installation again.

```
kubectl create namespace NAME
```

Error saving Docker credentials

To [set up vertical autoscaling](#), you must authenticate and fetch your Amazon EMR on EKS vertical autoscaling-related Docker images. When you do this, you might get an error like the following one if Docker isn't running:

```
aws ecr get-login-password \
  --region $REGION | docker login \
  --username AWS \
  --password-stdin $ACCOUNT_ID.dkr.ecr.$REGION.amazonaws.com

Error saving credentials: error storing credentials - err: exit status 1
out: 'Post "http://ipc/registry/credstore-updated": dial unix backend.sock: connect: no
such file or directory'
```

To resolve this issue, confirm that Docker is running or open Docker Desktop. Then, try to save your credentials again.

Troubleshooting Amazon EMR on EKS Spark operator

Refer to the following sections if you encounter problems with the Amazon EMR on EKS Spark operator. For more information including steps to complete the installation, see [Running Spark jobs with the Spark operator](#).

Error on Helm chart installation

If you followed the steps in [Install the Spark operator](#) and it returned a `INSTALLATION FAILED` error like the one below when you tried to install or verify the Helm chart, you might not have obtained the authentication tokens to the Amazon ECR repository for the operator.

To resolve this issue, repeat the step in [Install the Spark operator](#) to authenticate your Helm client to the Amazon ECR registry. Then, try the installation step again.

```
Error: INSTALLATION FAILED: Kubernetes cluster unreachable: the server has asked for
the client to provide credentials
```

UnsupportedFileSystemException: No FileSystem for scheme "s3"

You might encounter the following exception in thread "main":

```
org.apache.hadoop.fs.UnsupportedFileSystemException: No FileSystem for scheme "s3"
```

If this occurs, add the following exceptions to the SparkApplication spec:

```
hadoopConf:
```

```

# EMRFS filesystem
fs.s3.customAWSCredentialsProvider:
com.amazonaws.auth.WebIdentityTokenCredentialsProvider
fs.s3.impl: com.amazon.ws.emr.hadoop.fs.EmrFileSystem
fs.AbstractFileSystem.s3.impl: org.apache.hadoop.fs.s3.EMRFSDelegate
fs.s3.buffer.dir: /mnt/s3
fs.s3.getObject.initialSocketTimeoutMilliseconds: "2000"
mapreduce.fileoutputcommitter.algorithm.version.emr_internal_use_only.EmrFileSystem:
"2"
mapreduce.fileoutputcommitter.cleanup-
failures.ignored.emr_internal_use_only.EmrFileSystem: "true"
sparkConf:
  # Required for EMR Runtime
  spark.driver.extraClassPath: /usr/lib/hadoop-lzo/lib/*:/usr/lib/hadoop/hadoop-
aws.jar:/usr/share/aws/aws-java-sdk/*:/usr/share/aws/emr/emrfs/conf:/usr/share/aws/
emr/emrfs/lib/*:/usr/share/aws/emr/emrfs/auxlib/*:/usr/share/aws/emr/security/conf:/
usr/share/aws/emr/security/lib/*:/usr/share/aws/hmclient/lib/aws-glue-datacatalog-
spark-client.jar:/usr/share/java/Hive-JSON-Serde/hive-openx-serde.jar:/usr/share/aws/
sagemaker-spark-sdk/lib/sagemaker-spark-sdk.jar:/home/hadoop/extrajars/*
  spark.driver.extraLibraryPath: /usr/lib/hadoop/lib/native:/usr/lib/hadoop-lzo/lib/
native:/docker/usr/lib/hadoop/lib/native:/docker/usr/lib/hadoop-lzo/lib/native
  spark.executor.extraClassPath: /usr/lib/hadoop-lzo/lib/*:/usr/lib/hadoop/hadoop-
aws.jar:/usr/share/aws/aws-java-sdk/*:/usr/share/aws/emr/emrfs/conf:/usr/share/aws/
emr/emrfs/lib/*:/usr/share/aws/emr/emrfs/auxlib/*:/usr/share/aws/emr/security/conf:/
usr/share/aws/emr/security/lib/*:/usr/share/aws/hmclient/lib/aws-glue-datacatalog-
spark-client.jar:/usr/share/java/Hive-JSON-Serde/hive-openx-serde.jar:/usr/share/aws/
sagemaker-spark-sdk/lib/sagemaker-spark-sdk.jar:/home/hadoop/extrajars/*
  spark.executor.extraLibraryPath: /usr/lib/hadoop/lib/native:/usr/lib/hadoop-lzo/lib/
native:/docker/usr/lib/hadoop/lib/native:/docker/usr/lib/hadoop-lzo/lib/native

```

Amazon EMR on EKS service endpoints and quotas

The following are the service endpoints and service quotas for Amazon EMR on EKS. To connect programmatically to an AWS service, you use an endpoint. In addition to the standard AWS endpoints, some AWS services offer FIPS endpoints in selected Regions. For more information, see [AWS service endpoints](#). *Service quotas*, also referred to as *limits*, are the maximum number of service resources or operations for your AWS account. For more information, see [AWS service quotas](#).

Service endpoints

AWS Region name	Code	Endpoint	Protocol
US East (N. Virginia)	us-east-1	emr-containers.us-east-1.amazonaws.com	HTTPS
US East (Ohio)	us-east-2	emr-containers.us-east-2.amazonaws.com	HTTPS
US West (N. California)	us-west-1	emr-containers.us-west-1.amazonaws.com	HTTPS
US West (Oregon)	us-west-2	emr-containers.us-west-2.amazonaws.com	HTTPS
Asia Pacific (Tokyo)	ap-northeast-1	emr-containers.ap-northeast-1.amazonaws.com	HTTPS
Asia Pacific (Seoul)	ap-northeast-2	emr-containers.ap-northeast-2.amazonaws.com	HTTPS
Asia Pacific (Osaka)	ap-northeast-3	emr-containers.ap-northeast-3.amazonaws.com	HTTPS
Asia Pacific (Mumbai)	ap-south-1	emr-containers.ap-south-1.amazonaws.com	HTTPS

AWS Region name	Code	Endpoint	Protocol
Asia Pacific (Hyderabad)	ap-south-2	emr-containers.ap-south-2.amazonaws.com	HTTPS
Asia Pacific (Singapore)	ap-southeast-1	emr-containers.ap-southeast-1.amazonaws.com	HTTPS
Asia Pacific (Sydney)	ap-southeast-2	emr-containers.ap-southeast-2.amazonaws.com	HTTPS
Asia Pacific (Jakarta)	ap-southeast-3	emr-containers.ap-southeast-3.amazonaws.com	HTTPS
Asia Pacific (Hong Kong)	ap-east-1	emr-containers.ap-east-1.amazonaws.com	HTTPS
Africa (Cape Town)	af-south-1	emr-containers.af-south-1.amazonaws.com	HTTPS
Canada (Central)	ca-central-1	emr-containers.ca-central-1.amazonaws.com	HTTPS
China (Ningxia)	cn-northwest-1	emr-containers.cn-northwest-1.amazonaws.com.cn	HTTPS
China (Beijing)	cn-north-1	emr-containers.cn-north-1.amazonaws.com.cn	HTTPS
Europe (Frankfurt)	eu-central-1	emr-containers.eu-central-1.amazonaws.com	HTTPS
Europe (Zurich)	eu-central-2	emr-containers.eu-central-2.amazonaws.com	HTTPS
Europe (Ireland)	eu-west-1	emr-containers.eu-west-1.amazonaws.com	HTTPS

AWS Region name	Code	Endpoint	Protocol
Europe (London)	eu-west-2	emr-containers.eu-west-2.amazonaws.com	HTTPS
Europe (Paris)	eu-west-3	emr-containers.eu-west-3.amazonaws.com	HTTPS
Europe (Stockholm)	eu-north-1	emr-containers.eu-north-1.amazonaws.com	HTTPS
Europe (Milan)	eu-south-1	emr-containers.eu-south-1.amazonaws.com	HTTPS
Europe (Spain)	eu-south-2	emr-containers.eu-south-2.amazonaws.com	HTTPS
Israel (Tel Aviv)	il-central-1	emr-containers.il-central-1.amazonaws.com	HTTPS
South America (São Paulo)	sa-east-1	emr-containers.sa-east-1.amazonaws.com	HTTPS
Middle East (UAE)	me-central-1	emr-containers.me-central-1.amazonaws.com	HTTPS
Middle East (Bahrain)	me-south-1	emr-containers.me-south-1.amazonaws.com	HTTPS
AWS GovCloud (US-East)	us-gov-east-1	emr-containers.us-gov-east-1.amazonaws.com	HTTPS
AWS GovCloud (US-West)	us-gov-west-1	emr-containers.us-gov-west-1.amazonaws.com	HTTPS

Service quotas

Amazon EMR on EKS throttles the following API requests for each AWS account on a per-Region basis. For more information about how throttling is applied, see [API Request Throttling](#) in the

Amazon EC2 API Reference. You can request an increase to API throttling quotas for your AWS account, following the guide that follows.

API action	Bucket maximum capacity	Bucket refill rate (per second)
CancelJobRun	25	1
CreateJobTemplate	25	1
CreateManagedEndpoint	25	1
CreateSecurityConfiguration	25	1
CreateVirtualCluster	25	1
DeleteJobTemplate	25	1
DeleteManagedEndpoint	25	1
DeleteVirtualCluster	25	1
DescribeJobRun	100	20
DescribeJobTemplate	25	1
DescribeManagedEndpoint	100	5
DescribeSecurityConfiguration	25	1
DescribeVirtualCluster	100	5
GetManagedEndpointSessionCredentials	25	1
ListJobRuns	100	5
ListJobTemplates	25	1
ListManagedEndpoints	25	1

API action	Bucket maximum capacity	Bucket refill rate (per second)
ListSecurityConfigurations	25	1
ListVirtualClusters	100	5
StartJobRun	25	1
Throttle quota for all EMR on EKS API requests	200	20

When you create your AWS account, we set default quotas (also referred to as limits) on your AWS resources on a per-Region basis. If you attempt to exceed the quota for a resource, the request fails. If this happens, you can reduce your resource usage or request a quota increase.

The **Service Quotas** console is a central location where you can view and manage your quotas for AWS services, and request quota increases for many of the resources that you use. Use the quota information provided here to manage your AWS infrastructure. Plan to request any quota increases in advance of the time that you'll need them.

View quotas and request quota increases

You can view your current service quotas for each Region by using the **Service Quotas** console. For detailed instructions, see [Viewing Service quotas](#) in the *Service Quotas User Guide*.

You can request a quota increase from either the AWS Management Console or using the AWS CLI. The steps are detailed in [Requesting a quota increase](#).

Amazon EMR on EKS releases

An Amazon EMR release is a set of open-source applications from the big data ecosystem. Each release comprises different big data applications, components, and features that you select to have Amazon EMR on EKS deploy and configure when you run your job.

Beginning with Amazon EMR releases 5.32.0 and 6.2.0, you can deploy Amazon EMR on EKS. This deployment option is not available with earlier Amazon EMR release versions. You must specify a supported release version when you submit your job.

Amazon EMR on EKS uses the following form of release label: `emr-x.x.x-latest` or `emr-x.x.x-yyyyymmdd` with a specific release date. For example, `emr-7.13.0-latest` or `emr-7.13.0-20210129`. When you use the `-latest` suffix, you ensure that your Amazon EMR version always includes the latest security updates.

Note

For a comparison between Amazon EMR on EKS and Amazon EMR running on EC2, see the [Amazon EMR FAQs](#) on the AWS website.

Topics

- [AWS runtime for Apache Spark \(emr-spark-8.0.0\) on EKS](#)
- [Amazon EMR on EKS 7.13.0 releases](#)
- [Amazon EMR on EKS 7.12.0 releases](#)
- [Amazon EMR on EKS 7.11.0 releases](#)
- [Amazon EMR on EKS 7.10.0 releases](#)
- [Amazon EMR on EKS 7.9.0 releases](#)
- [Amazon EMR on EKS 7.8.0 releases](#)
- [Amazon EMR on EKS 7.7.0 releases](#)
- [Amazon EMR on EKS 7.6.0 releases](#)
- [Amazon EMR on EKS 7.5.0 releases](#)
- [Amazon EMR on EKS 7.4.0 releases](#)

- [Amazon EMR on EKS 7.3.0 releases](#)
- [Amazon EMR on EKS 7.2.0 releases](#)
- [Amazon EMR on EKS 7.1.0 releases](#)
- [Amazon EMR on EKS 7.0.0 releases](#)
- [Amazon EMR on EKS 6.15.0 releases](#)
- [Amazon EMR on EKS 6.14.0 releases](#)
- [Amazon EMR on EKS 6.13.0 releases](#)
- [Amazon EMR on EKS 6.12.0 releases](#)
- [Amazon EMR on EKS 6.11.0 releases](#)
- [Amazon EMR on EKS 6.10.0 releases](#)
- [Amazon EMR on EKS 6.9.0 releases](#)
- [Amazon EMR on EKS 6.8.0 releases](#)
- [Amazon EMR on EKS 6.7.0 releases](#)
- [Amazon EMR on EKS 6.6.0 releases](#)
- [Amazon EMR on EKS 6.5.0 releases](#)
- [Amazon EMR on EKS 6.4.0 releases](#)
- [Amazon EMR on EKS 6.3.0 releases](#)
- [Amazon EMR on EKS 6.2.0 releases](#)
- [Amazon EMR on EKS 5.36.0 releases](#)
- [Amazon EMR on EKS 5.35.0 releases](#)
- [Amazon EMR on EKS 5.34.0 releases](#)
- [Amazon EMR on EKS 5.33.0 releases](#)
- [Amazon EMR on EKS 5.32.0 releases](#)

AWS runtime for Apache Spark (emr-spark-8.0.0) on EKS

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR Spark 8.0.0 release in general, see [AWS runtime for Apache Spark \(emr-spark-8.0.0\)](#) in the *Amazon EMR Release Guide*.

AWS runtime for Apache Spark (emr-spark-8.0.0) on EKS

The following `emr-spark-8.0.0` releases are available for AWS runtime for Apache Spark on EKS. Select a specific **emr-spark-8.0.0-XXXX** release to view more details such as the related container image tag.

Spark releases

The following `emr-spark-8.0.0` releases are available for AWS runtime for Apache Spark on EKS when you run Spark applications. Select a `spark/` release to view details such as the release label and container image tag. The `notebook-spark/`, `notebook-python/`, and `livy/` variants use the same release label format (for example, `emr-spark-8.0.0-latest`) with the corresponding container image prefix.

- [emr-spark-8.0.0-latest](#)
- [emr-spark-8.0.0-20260421](#)
- `notebook-spark/emr-spark-8.0.0-latest`
- `notebook-spark/emr-spark-8.0.0-20260421`
- `notebook-python/emr-spark-8.0.0-latest`
- `notebook-python/emr-spark-8.0.0-20260421`
- `livy/emr-spark-8.0.0-latest`
- `livy/emr-spark-8.0.0-20260421`

Release notes

Release notes for AWS runtime for Apache Spark (`emr-spark-8.0.0`) on EKS:

- **Supported applications** - AWS SDK for Java 2.41.32, Apache Spark 4.0.2-amzn-0, Apache Hudi 1.1.0-amzn-0, Apache Iceberg 1.10.1-amzn-0, Delta Lake 4.0.0-amzn-1-spark
- **Supported components** - `emr-ddb`, `emr-goodies`, `hadoop-client`, `hudi`, `hudi-spark`, `iceberg`, `spark-kubernetes`.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the <code>core-site.xml</code> Hadoop file.
spark-metrics	Change values in the <code>metrics.properties</code> Spark file.
spark-defaults	Change values in the <code>spark-defaults.conf</code> Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the <code>hive-site.xml</code> Spark file.
spark-log4j2	Change values in the <code>log4j2.properties</code> Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Changes and features

The following features are included with the `emr-spark-8.0.0` release of AWS runtime for Apache Spark on EKS:

- **Apache Spark 4.0.2 GA** – First production-ready release of Spark 4.x on Amazon EMR on EKS, featuring ANSI SQL mode, SQL PIPE syntax, VARIANT data type, SQL scripting, and streaming enhancements.
- **Python 3.11 default** – Python 3.11 is the default for PySpark and Spark workloads. Python 3.12 and 3.13 are also available.

`emr-spark-8.0.0-latest`

Release label

Use `emr-spark-8.0.0-latest` as the `--release-label` value in the [StartJobRun](#) API.

Release notes

`emr-spark-8.0.0-latest` currently points to `emr-spark-8.0.0-20260421`.

Regions

You can use `emr-spark-8.0.0-latest` in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag

`spark/emr-spark-8.0.0:latest`

`emr-spark-8.0.0-20260421`

Release label

Use `emr-spark-8.0.0-20260421` as the `--release-label` value in the [StartJobRun](#) API.

Release notes

`emr-spark-8.0.0-20260421` was released in April 2026. This is the initial release of AWS runtime for Apache Spark 8.0.0.

Regions

You can use `emr-spark-8.0.0-20260421` in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag

`spark/emr-spark-8.0.0:20260421`

Amazon EMR on EKS 7.13.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.13.0 release in general, see [Amazon EMR 7.13.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.13 releases

The following Amazon EMR 7.13.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.13.0-XXXX** release to view more details such as the related container image tag.

Flink releases

The following Amazon EMR 7.13.0 releases are available for Amazon EMR on EKS when you run Flink applications.

- [emr-7.13.0-flink-latest](#)
- [emr-7.13.0-flink-20260410](#)

Spark releases

The following Amazon EMR 7.13.0 releases are available for Amazon EMR on EKS when you run Spark applications.

- [emr-7.13.0-latest](#)
- [emr-7.13.0-20260410](#)
- `emr-7.13.0-spark-rapids-latest`
- `emr-7.13.0-spark-rapids-20260410`

- `emr-7.13.0-java11-latest`
- `emr-7.13.0-java11-20260410`
- `emr-7.13.0-java8-latest`
- `emr-7.13.0-java8-20260410`
- `emr-7.13.0-spark-rapids-java8-latest`
- `emr-7.13.0-spark-rapids-java8-20260410`
- `notebook-spark/emr-7.13.0-latest`
- `notebook-spark/emr-7.13.0-20260410`
- `notebook-spark/emr-7.13.0-spark-rapids-latest`
- `notebook-spark/emr-7.13.0-spark-rapids-20260410`
- `notebook-spark/emr-7.13.0-java11-latest`
- `notebook-spark/emr-7.13.0-java11-20260410`
- `notebook-spark/emr-7.13.0-java8-latest`
- `notebook-spark/emr-7.13.0-java8-20260410`
- `notebook-spark/emr-7.13.0-spark-rapids-java8-latest`
- `notebook-spark/emr-7.13.0-spark-rapids-java8-20260410`
- `notebook-python/emr-7.13.0-latest`
- `notebook-python/emr-7.13.0-20260410`
- `notebook-python/emr-7.13.0-spark-rapids-latest`
- `notebook-python/emr-7.13.0-spark-rapids-20260410`
- `notebook-python/emr-7.13.0-java11-latest`
- `notebook-python/emr-7.13.0-java11-20260410`
- `notebook-python/emr-7.13.0-java8-latest`
- `notebook-python/emr-7.13.0-java8-20260410`
- `notebook-python/emr-7.13.0-spark-rapids-java8-latest`
- `notebook-python/emr-7.13.0-spark-rapids-java8-20260410`
- `livy/emr-7.13.0-latest`
- `livy/emr-7.13.0-20260410`
- `livy/emr-7.13.0-java11-latest`
- `livy/emr-7.13.0-java11-20260410`

- `livy/emr-7.13.0-java8-latest`
- `livy/emr-7.13.0-java8-20260410`

Release notes

Release notes for Amazon EMR on EKS 7.13.0:

- **Supported applications** - AWS SDK for Java 2.42.12 and 1.12.797, Apache Spark 3.5.6-amzn-2, Apache Hudi 1.0.2-amzn-2, Apache Iceberg 1.10.0-amzn-1, Delta 3.3.2-amzn-2, Apache Spark RAPIDS 25.04.0-amzn-0, Apache Flink 1.20.0-amzn-7
- **Supported components** - `emr-ddb`, `emr-goodies`, `emr-s3-select`, `emrfs`, `hadoop-client`, `hudi`, `hudi-spark`, `iceberg`, `spark-kubernetes`.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
<code>core-site</code>	Change values in the <code>core-site.xml</code> Hadoop file.
<code>emrfs-site</code>	Change EMRFS settings.
<code>spark-metrics</code>	Change values in the <code>metrics.properties</code> Spark file.
<code>spark-defaults</code>	Change values in the <code>spark-defaults.conf</code> Spark file.
<code>spark-env</code>	Change values in the Spark environment.
<code>spark-hive-site</code>	Change values in the <code>hive-site.xml</code> Spark file.
<code>spark-log4j2</code>	Change values in the <code>log4j2.properties</code> Spark file.
<code>emr-job-submitter</code>	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Changes and features

The following features are included with the 7.13.0 release of Amazon EMR on EKS:

- **Python 3.11 default for PySpark and Spark workloads** – Python 3.11 is now the default Python version for PySpark and Spark workloads. Python 3.9 remains the default for all other applications. Both Python 3.9 and 3.11 are included in the release.

emr-7.13.0-latest

Release notes: `emr-7.13.0-latest` currently points to `emr-7.13.0-20260410`.

Regions: `emr-7.13.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.13.0:latest`

emr-7.13.0-20260410

Release notes: `emr-7.13.0-20260410` was released in April 2026. This is the initial release of Amazon EMR 7.13.0 (Spark).

Regions: `emr-7.13.0-20260410` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.13.0-20260410`

emr-7.13.0 common vulnerabilities and exposures

The following table lists all CVEs that don't impact EMR on EKS job runs that run on recommended configurations of `emr-7.13.0`. Amazon EMR is dependent on upstream open-source for availability of fixes and will provide the latest stable version as part of the Core Engine components within 90 days of Amazon EMR verifying the fixes.

This table was updated on August 21, 2026.

CVE ID	Severity	CVE details URL
CVE-2026-2332	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2026-2332

emr-7.13.0-flink-latest

Release notes: `emr-7.13.0-flink-latest` currently points to `emr-7.13.0-flink-20260410`

Regions: `emr-7.13.0-flink-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.13.0-flink:latest`

emr-7.13.0-flink-20260410

Release notes: `emr-7.13.0-flink-20260410` was released in April 2026. This is the initial release of Amazon EMR 7.13.0 (Flink).

Regions: `emr-7.13.0-flink-20260410` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.13.0-flink:20260410`

Amazon EMR on EKS 7.12.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.12.0 release in general, see [Amazon EMR 7.12.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.12 releases

The following Amazon EMR 7.12.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.12.0-XXXX** release to view more details such as the related container image tag.

Flink releases

The following Amazon EMR 7.12.0 releases are available for Amazon EMR on EKS when you run Flink applications.

- [emr-7.12.0-flink-latest](#)
- [emr-7.12.0-flink-20251111](#)

Spark releases

The following Amazon EMR 7.12.0 releases are available for Amazon EMR on EKS when you run Spark applications.

- [emr-7.12.0-latest](#)
- [emr-7.12.0-20251111](#)
- emr-7.12.0-spark-rapids-latest
- emr-7.12.0-spark-rapids-20251111
- emr-7.12.0-java11-latest
- emr-7.12.0-java11-20251111
- emr-7.12.0-java8-latest
- emr-7.12.0-java8-20251111
- emr-7.12.0-spark-rapids-java8-latest
- emr-7.12.0-spark-rapids-java8-20251111

- notebook-spark/emr-7.12.0-latest
- notebook-spark/emr-7.12.0-20251111
- notebook-spark/emr-7.12.0-spark-rapids-latest
- notebook-spark/emr-7.12.0-spark-rapids-20251111
- notebook-spark/emr-7.12.0-java11-latest
- notebook-spark/emr-7.12.0-java11-20251111
- notebook-spark/emr-7.12.0-java8-latest
- notebook-spark/emr-7.12.0-java8-20251111
- notebook-spark/emr-7.12.0-spark-rapids-java8-latest
- notebook-spark/emr-7.12.0-spark-rapids-java8-20251111
- notebook-python/emr-7.12.0-latest
- notebook-python/emr-7.12.0-20251111
- notebook-python/emr-7.12.0-spark-rapids-latest
- notebook-python/emr-7.12.0-spark-rapids-20251111
- notebook-python/emr-7.12.0-java11-latest
- notebook-python/emr-7.12.0-java11-20251111
- notebook-python/emr-7.12.0-java8-latest
- notebook-python/emr-7.12.0-java8-20251111
- notebook-python/emr-7.12.0-spark-rapids-java8-latest
- notebook-python/emr-7.12.0-spark-rapids-java8-20251111
- livy/emr-7.12.0-latest
- livy/emr-7.12.0-20251111
- livy/emr-7.12.0-java11-latest
- livy/emr-7.12.0-java11-20251111
- livy/emr-7.12.0-java8-latest
- livy/emr-7.12.0-java8-20251111

Release notes

Release notes for Amazon EMR on EKS 7.12.0:

- **Supported applications** - AWS SDK for Java 2.35.5 and 1.12.792, Apache Spark 3.5.6-amzn-1, Apache Hudi 1.0.2-amzn-1, Apache Iceberg 1.10.0-amzn-0, Delta 3.3.2-amzn-1, Apache Spark RAPIDS 25.04.0-amzn-0, Apache Flink 1.20.0-amzn-6
- **Supported components** - emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the core-site.xml Hadoop file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the metrics.properties Spark file.
spark-defaults	Change values in the spark-defaults.conf Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the hive-site.xml Spark file.
spark-log4j2	Change values in the log4j2.properties Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Changes and features

The following features are included with the 7.12.0 release of Amazon EMR on EKS:

- **Iceberg Materialized Views** – Starting EMR 7.12.0, EMR Spark supports creation and management of Iceberg Materialized Views (MV).
- **Hudi Full Table Access** – Starting EMR 7.12.0, EMR now supports Full Table Access (FTA) control for Apache Hudi in Apache Spark based on your policies defined in Lake Formation. This feature enables read and write operations from your Amazon EMR Spark jobs on Lake Formation registered tables when the job role has full table access.
- **Iceberg version upgrade** – EMR 7.12.0 supports Apache Iceberg version 1.10.
- **Logging for Livy interactive workloads** – Starting EMR 7.12.0, EMR supports extensive logging for key system components to improve troubleshooting for Livy Spark job failures. This feature will provide EMR service with access to additional Livy and SecretAgent logs to simplify troubleshooting.

emr-7.12.0-latest

Release notes: `emr-7.12.0-latest` currently points to `emr-7.12.0-20251111`.

Regions: `emr-7.12.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.12.0:latest`

emr-7.12.0-20251111

Release notes: `emr-7.12.0-20251111` was released in November 2025. This is the initial release of Amazon EMR 7.12.0 (Spark).

Regions: `emr-emr-7.12.0-20251111` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.12.0-20251111`

emr-7.12.0 common vulnerabilities and exposures

The following table lists all CVEs that don't impact EMR on EKS job runs that run on recommended configurations of `emr-7.12.0`. Amazon EMR is dependent on upstream open-source for availability of fixes and will provide the latest stable version as part of the Core Engine components within 90 days of Amazon EMR verifying the fixes.

This table was updated on August 21, 2026.

CVE ID	Severity	CVE details URL
CVE-2024-51504	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2024-51504
CVE-2025-30065	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2025-30065
CVE-2026-2332	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2026-2332
CVE-2024-7254	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-7254
CVE-2025-48734	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-48734
CVE-2025-52999	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-52999

CVE ID	Severity	CVE details URL
CVE-2026-5450	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2026-5450
CVE-2025-54920	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-54920

emr-7.12.0-flink-latest

Release notes: `emr-7.12.0-flink-latest` currently points to `emr-7.12.0-flink-20251111`

Regions: `emr-7.12.0-flink-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.12.0-flink:latest`

emr-7.12.0-flink-20251111

Release notes: `7.12.0-flink-20251111` was released in November 2025. This is the initial release of Amazon EMR 7.12.0 (Flink).

Regions: `emr-7.12.0-flink-20251111` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.12.0-flink:20251111`

Amazon EMR on EKS 7.11.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.11.0 release in general, see [Amazon EMR 7.11.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.11 releases

The following Amazon EMR 7.11.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.11.0-XXXX** release to view more details such as the related container image tag.

Flink releases

The following Amazon EMR 7.11.0 releases are available for Amazon EMR on EKS when you run Flink applications.

- [emr-7.11.0-flink-latest](#)
- [emr-7.11.0-flink-20251020](#)

Spark releases

The following Amazon EMR 7.11.0 releases are available for Amazon EMR on EKS when you run Spark applications.

- [emr-7.11.0-latest](#)
- [emr-7.11.0-20251020](#)
- emr-7.11.0-spark-rapids-latest
- emr-7.11.0-spark-rapids-20251020
- emr-7.11.0-java11-latest
- emr-7.11.0-java11-20251020
- emr-7.11.0-java8-latest
- emr-7.11.0-java8-20251020
- emr-7.11.0-spark-rapids-java8-latest
- emr-7.11.0-spark-rapids-java8-20251020
- notebook-spark/emr-7.11.0-latest
- notebook-spark/emr-7.11.0-20251020
- notebook-spark/emr-7.11.0-spark-rapids-latest
- notebook-spark/emr-7.11.0-spark-rapids-20251020
- notebook-spark/emr-7.11.0-java11-latest
- notebook-spark/emr-7.11.0-java11-20251020
- notebook-spark/emr-7.11.0-java8-latest
- notebook-spark/emr-7.11.0-java8-20251020
- notebook-spark/emr-7.11.0-spark-rapids-java8-latest
- notebook-spark/emr-7.11.0-spark-rapids-java8-20251020
- notebook-python/emr-7.11.0-latest

- notebook-python/emr-7.11.0-20251020
- notebook-python/emr-7.11.0-spark-rapids-latest
- notebook-python/emr-7.11.0-spark-rapids-20251020
- notebook-python/emr-7.11.0-java11-latest
- notebook-python/emr-7.11.0-java11-20251020
- notebook-python/emr-7.11.0-java8-latest
- notebook-python/emr-7.11.0-java8-20251020
- notebook-python/emr-7.11.0-spark-rapids-java8-latest
- notebook-python/emr-7.11.0-spark-rapids-java8-20251020
- livy/emr-7.11.0-latest
- livy/emr-7.11.0-20251020
- livy/emr-7.11.0-java11-latest
- livy/emr-7.11.0-java11-20251020
- livy/emr-7.11.0-java8-latest
- livy/emr-7.11.0-java8-20251020

Release notes

Release notes for Amazon EMR on EKS 7.11.0:

- **Supported applications** - AWS SDK for Java 2.35.5 and 1.12.792, Apache Spark 3.5.6-amzn-0, Apache Hudi 1.0.2-amzn-0, Apache Iceberg 1.9.1-amzn-0, Delta 3.3.2-amzn-0, Apache Spark RAPIDS 25.04.0-amzn-0, Apache Flink 1.20.0-amzn-5
- **Supported components** - emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the core-site.xml Hadoop file.

Classifications	Descriptions
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the <code>metrics.properties</code> Spark file.
spark-defaults	Change values in the <code>spark-defaults.conf</code> Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the <code>hive-site.xml</code> Spark file.
spark-log4j2	Change values in the <code>log4j2.properties</code> Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Changes and features

The following changes are included with the 7.11.0 release of Amazon EMR on EKS:

Amazon EMR on EKS now supports integration with SageMaker Unified Studio through a managed Livy solution including:

- **Managed Sessions:** New interactive session resource type that provides a custom HTTPS endpoint for running Spark sessions on your EKS cluster through SageMaker Unified Studio
- **Lake Formation Integration:** Supports data access control with two modes a) Fine-Grained Access Control b) Full Table Access (Compatibility Mode)
- **Identity Management:** Flexible authentication options a) IAM role-based access control b) based access control.
- **User Background Sessions with integration:** Supports long-running Spark workloads to continue running even after users log off from SageMaker Unified Studio, supporting sessions up to 90 days

emr-7.11.0-latest

Release notes: `emr-7.11.0-latest` currently points to `emr-7.11.0-20251020`.

Regions: `emr-7.11.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.11.0:latest`

emr-7.11.0-20251020

Release notes: `emr-7.11.0-20251020` was released in November 2025. This is the initial release of Amazon EMR 7.11.0 (Spark).

Regions: `emr-emr-7.11.0-20251020` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.11.0-20251020`

emr-7.11.0-flink-latest

Release notes: `emr-7.11.0-flink-latest` currently points to `emr-7.11.0-flink-20251020`

Regions: `emr-7.11.0-flink-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.11.0-flink:latest`

emr-7.11.0-flink-20251020

Release notes: `7.11.0-flink-20251020` was released in November 2025. This is the initial release of Amazon EMR 7.11.0 (Flink).

Regions: `emr-7.11.0-flink-20251020` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.11.0-flink:20251020`

Amazon EMR on EKS 7.10.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.10.0 release in general, see [Amazon EMR 7.10.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.10 releases

The following Amazon EMR 7.10.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.10.0-XXXX** release to view more details such as the related container image tag.

Flink releases

The following Amazon EMR 7.10.0 releases are available for Amazon EMR on EKS when you run Flink applications.

- [emr-7.10.0-flink-latest](#)
- [emr-7.10.0-flink-20250801](#)

Spark releases

The following Amazon EMR 7.10.0 releases are available for Amazon EMR on EKS when you run Spark applications.

- [emr-7.10.0-latest](#)
- [emr-7.10.0-20250801](#)
- emr-7.10.0-spark-rapids-latest
- emr-7.10.0-spark-rapids-20250801
- emr-7.10.0-java11-latest
- emr-7.10.0-java11-20250801
- emr-7.10.0-java8-latest
- emr-7.10.0-java8-20250801
- emr-7.10.0-spark-rapids-java8-latest
- emr-7.10.0-spark-rapids-java8-20250801
- notebook-spark/emr-7.10.0-latest
- notebook-spark/emr-7.10.0-20250801
- notebook-spark/emr-7.10.0-spark-rapids-latest
- notebook-spark/emr-7.10.0-spark-rapids-20250801
- notebook-spark/emr-7.10.0-java11-latest
- notebook-spark/emr-7.10.0-java11-20250801
- notebook-spark/emr-7.10.0-java8-latest
- notebook-spark/emr-7.10.0-java8-20250801
- notebook-spark/emr-7.10.0-spark-rapids-java8-latest
- notebook-spark/emr-7.10.0-spark-rapids-java8-20250801
- notebook-python/emr-7.10.0-latest
- notebook-python/emr-7.10.0-20250801
- notebook-python/emr-7.10.0-spark-rapids-latest
- notebook-python/emr-7.10.0-spark-rapids-20250801
- notebook-python/emr-7.10.0-java11-latest
- notebook-python/emr-7.10.0-java11-20250801
- notebook-python/emr-7.10.0-java8-latest
- notebook-python/emr-7.10.0-java8-20250801

- notebook-python/emr-7.10.0-spark-rapids-java8-latest
- notebook-python/emr-7.10.0-spark-rapids-java8-20250801
- livy/emr-7.10.0-latest
- livy/emr-7.10.0-20250801
- livy/emr-7.10.0-java11-latest
- livy/emr-7.10.0-java11-20250801
- livy/emr-7.10.0-java8-latest
- livy/emr-7.10.0-java8-20250801

Release notes

Release notes for Amazon EMR on EKS 7.10.0:

- **Supported applications** - AWS SDK for Java 2.31.48 and 1.12.782, Apache Spark 3.5.5-amzn-1, Apache Hudi 0.15.0-amzn-7, Apache Iceberg 1.8.1-amzn-0, Delta 3.3.0-amzn-2, Apache Spark RAPIDS 25.04.0-amzn-0, Apache Flink 1.20.0-amzn-4, Flink Kubernetes Operator 1.10.0-amzn-4
- **Supported components** - emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the core-site.xml Hadoop file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the metrics.properties Spark file.
spark-defaults	Change values in the spark-defaults.conf Spark file.
spark-env	Change values in the Spark environment.

Classifications	Descriptions
spark-hive-site	Change values in the <code>hive-site.xml</code> Spark file.
spark-log4j2	Change values in the <code>log4j2.properties</code> Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Changes and features

The following features are included with the 7.10.0 release of Amazon EMR on EKS:

- **S3A filesystem** – Starting from the 7.10.0 release, the S3A filesystem has replaced EMRFS as the default EMR S3 connector. For more information, see [EMR File System \(EMRFS\)](#).

emr-7.10.0-latest

Release notes: `emr-7.10.0-latest` currently points to `emr-7.10.0-20250801`.

Regions: `emr-7.10.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.10.0:latest`

emr-7.10.0-20250801

Release notes: `emr-7.10.0-20250801` was released in February 2025. This is the initial release of Amazon EMR 7.10.0 (Spark).

Regions: `emr-emr-7.10.0-20250801` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.10.0-20250801`

emr-7.10.0 common vulnerabilities and exposures

The following table lists all CVEs that don't impact EMR on EKS job runs that run on recommended configurations of `emr-7.10.0`. Amazon EMR is dependent on upstream open-source for availability of fixes and will provide the latest stable version as part of the Core Engine components within 90 days of Amazon EMR verifying the fixes.

This table was updated on August 21, 2026.

CVE ID	Severity	CVE details URL
CVE-2026-25547	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2026-25547
CVE-2026-5450	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2026-5450
CVE-2007-3641	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2007-3641
CVE-2019-10172	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2019-10172
CVE-2021-33503	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2021-33503

CVE ID	Severity	CVE details URL
CVE-2021-38604	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2021-38604
CVE-2021-43396	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2021-43396
CVE-2021-46877	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2021-46877
CVE-2022-46337	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2022-46337
CVE-2023-29483	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2023-29483
CVE-2023-31484	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2023-31484
CVE-2023-45288	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2023-45288
CVE-2023-47100	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2023-47100
CVE-2023-49083	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2023-49083
CVE-2023-4911	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2023-4911
CVE-2023-5072	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2023-5072
CVE-2024-24790	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2024-24790
CVE-2024-25062	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2024-25062

CVE ID	Severity	CVE details URL
CVE-2024-34459	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2024-34459
CVE-2024-3651	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2024-3651
CVE-2024-53899	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2024-53899
CVE-2024-56171	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2024-56171
CVE-2025-0665	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-0665
CVE-2025-15467	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-15467
CVE-2025-22871	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-22871
CVE-2025-24970	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-24970
CVE-2025-27113	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-27113
CVE-2025-30065	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-30065
CVE-2025-32414	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-32414
CVE-2025-32415	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-32415
CVE-2025-46762	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-46762

CVE ID	Severity	CVE details URL
CVE-2025-47273	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-47273
CVE-2025-4802	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-4802
CVE-2025-53000	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-53000
CVE-2025-54920	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-54920
CVE-2025-6021	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-6021
CVE-2025-64756	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-64756
CVE-2025-68121	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-68121
CVE-2025-9086	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2025-9086
CVE-2026-0861	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2026-0861
CVE-2026-0915	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2026-0915
CVE-2026-24842	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2026-24842
CVE-2026-26960	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2026-26960
CVE-2026-26996	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2026-26996

CVE ID	Severity	CVE details URL
CVE-2026-27903	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2026-27903
CVE-2026-27904	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2026-27904
CVE-2026-30922	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2026-30922
CVE-2026-32597	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2026-32597
CVE-2026-33870	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2026-33870
CVE-2026-4176	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2026-4176
CVE-2026-5928	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2026-5928

emr-7.10.0-flink-latest

Release notes: `emr-7.10.0-flink-latest` currently points to `emr-7.10.0-flink-20250801`

Regions: `emr-7.10.0-flink-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.10.0-flink:latest`

emr-7.10.0-flink-20250801

Release notes: `7.10.0-flink-20250801` was released in February 2025. This is the initial release of Amazon EMR 7.10.0 (Flink).

Regions: `emr-7.10.0-flink-20250801` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.10.0-flink:20250801`

Amazon EMR on EKS 7.9.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.9.0 release in general, see [Amazon EMR 7.9.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.9 releases

The following Amazon EMR 7.9.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.9.0-XXXX** release to view more details such as the related container image tag.

Flink releases

The following Amazon EMR 7.9.0 releases are available for Amazon EMR on EKS when you run Flink applications.

- [emr-7.9.0-flink-latest](#)
- [emr-7.9.0-flink-20250425](#)

Spark releases

The following Amazon EMR 7.9.0 releases are available for Amazon EMR on EKS when you run Spark applications.

- [emr-7.9.0-latest](#)
- [emr-7.9.0-20250425](#)
- `emr-7.9.0-spark-rapids-latest`
- `emr-7.9.0-spark-rapids-20250425`
- `emr-7.9.0-java11-latest`
- `emr-7.9.0-java11-20250425`
- `emr-7.9.0-java8-latest`
- `emr-7.9.0-java8-20250425`
- `emr-7.9.0-spark-rapids-java8-latest`

- `emr-7.9.0-spark-rapids-java8-20250425`
- `notebook-spark/emr-7.9.0-latest`
- `notebook-spark/emr-7.9.0-20250425`
- `notebook-spark/emr-7.9.0-spark-rapids-latest`
- `notebook-spark/emr-7.9.0-spark-rapids-20250425`
- `notebook-spark/emr-7.9.0-java11-latest`
- `notebook-spark/emr-7.9.0-java11-20250425`
- `notebook-spark/emr-7.9.0-java8-latest`
- `notebook-spark/emr-7.9.0-java8-20250425`
- `notebook-spark/emr-7.9.0-spark-rapids-java8-latest`
- `notebook-spark/emr-7.9.0-spark-rapids-java8-20250425`
- `notebook-python/emr-7.9.0-latest`
- `notebook-python/emr-7.9.0-20250425`
- `notebook-python/emr-7.9.0-spark-rapids-latest`
- `notebook-python/emr-7.9.0-spark-rapids-20250425`
- `notebook-python/emr-7.9.0-java11-latest`
- `notebook-python/emr-7.9.0-java11-20250425`
- `notebook-python/emr-7.9.0-java8-latest`
- `notebook-python/emr-7.9.0-java8-20250425`
- `notebook-python/emr-7.9.0-spark-rapids-java8-latest`
- `notebook-python/emr-7.9.0-spark-rapids-java8-20250425`
- `livy/emr-7.9.0-latest`
- `livy/emr-7.9.0-20250425`
- `livy/emr-7.9.0-java11-latest`
- `livy/emr-7.9.0-java11-20250425`
- `livy/emr-7.9.0-java8-latest`
- `livy/emr-7.9.0-java8-20250425`

Release notes

Release notes for Amazon EMR on EKS 7.9.0

- **Supported applications** - AWS SDK for Java 2.31.16 and 1.12.782, Apache Spark 3.5.5, Apache Hudi 0.15.0-amzn-6, Apache Iceberg 1.7.1-amzn-2, Delta 3.3.0-amzn-1, Apache Spark RAPIDS 25.02.1-amzn-0, Jupyter Enterprise Gateway 2.6.0, Apache Flink 1.20.0-amzn-3, Flink Operator 1.10.0-amzn-3
- **Supported components** - emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the core-site.xml Hadoop file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the metrics.properties Spark file.
spark-defaults	Change values in the spark-defaults.conf Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the hive-site.xml Spark file.
spark-log4j2	Change values in the log4j2.properties Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Changes

The following changes are included with the 7.9.0 release of Amazon EMR on EKS:

- No changes for the release.

emr-7.9.0-latest

Release notes: `emr-7.9.0-latest` currently points to `emr-7.9.0-20250425`.

Regions: `emr-7.9.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.9.0:latest`

emr-7.9.0-20250425

Release notes: `emr-7.9.0-20250425` was released in February 2025. This is the initial release of Amazon EMR 7.9.0 (Spark).

Regions: `emr-emr-7.9.0-20250425` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.9.0-20250425`

emr-7.9.0 common vulnerabilities and exposures

The following table lists all CVEs that don't impact EMR on EKS job runs that run on recommended configurations of emr-7.9.0. Amazon EMR is dependent on upstream open-source for availability of fixes and will provide the latest stable version as part of the Core Engine components within 90 days of Amazon EMR verifying the fixes.

This table was updated on August 21, 2026.

CVE ID	Severity	CVE details URL
CVE-2019-20478	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2019-20478
CVE-2023-0687	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2023-0687
CVE-2023-47100	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2023-47100
CVE-2024-51504	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2024-51504
CVE-2025-30065	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2025-30065
CVE-2025-46762	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2025-46762
CVE-2019-0205	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2019-0205
CVE-2020-13949	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2020-13949
CVE-2020-28052	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2020-28052
CVE-2021-33503	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2021-33503

CVE ID	Severity	CVE details URL
CVE-2021-38604	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2021-38604
CVE-2021-3998	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2021-3998
CVE-2021-43396	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2021-43396
CVE-2022-46751	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2022-46751
CVE-2023-0286	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-0286
CVE-2023-29483	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-29483
CVE-2023-2976	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-2976
CVE-2023-31484	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-31484
CVE-2023-31486	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-31486
CVE-2023-43804	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-43804
CVE-2023-47038	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-47038
CVE-2023-49083	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-49083
CVE-2023-4911	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-4911

CVE ID	Severity	CVE details URL
CVE-2023-5072	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-5072
CVE-2023-50782	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-50782
CVE-2023-5156	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-5156
CVE-2024-13009	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-13009
CVE-2024-21634	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-21634
CVE-2024-2398	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-2398
CVE-2024-26130	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-26130
CVE-2024-33599	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-33599
CVE-2024-33601	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-33601
CVE-2024-33602	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-33602
CVE-2024-36114	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-36114
CVE-2024-3651	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-3651
CVE-2024-53899	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-53899

CVE ID	Severity	CVE details URL
CVE-2024-6119	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-6119
CVE-2024-6345	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-6345
CVE-2024-6365	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-6365
CVE-2024-7254	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-7254
CVE-2024-9823	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-9823
CVE-2025-24294	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-24294
CVE-2025-24970	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-24970
CVE-2025-27820	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-27820
CVE-2025-30749	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-30749
CVE-2025-47273	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-47273
CVE-2025-4802	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-4802
CVE-2025-48734	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-48734
CVE-2025-50059	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-50059

CVE ID	Severity	CVE details URL
CVE-2025-50106	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-50106

emr-7.9.0-flink-latest

Release notes: `emr-7.9.0-flink-latest` currently points to `emr-7.9.0-flink-20250425`

Regions: `emr-7.9.0-flink-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.9.0-flink:latest`

emr-7.9.0-flink-20250425

Release notes: `7.9.0-flink-20250425` was released in February 2025. This is the initial release of Amazon EMR 7.9.0 (Flink).

Regions: `emr-7.9.0-flink-20250425` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.9.0-flink:20250425`

Amazon EMR on EKS 7.8.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.8.0 release in general, see [Amazon EMR 7.8.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.8 releases

The following Amazon EMR 7.8.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.8.0-XXXX** release to view more details such as the related container image tag.

Flink releases

The following Amazon EMR 7.8.0 releases are available for Amazon EMR on EKS when you run Flink applications.

- [emr-7.8.0-flink-latest](#)
- [emr-7.8.0-flink-20250228](#)

Spark releases

The following Amazon EMR 7.8.0 releases are available for Amazon EMR on EKS when you run Spark applications.

- [emr-7.8.0-latest](#)
- [emr-7.8.0-20250228](#)
- emr-7.8.0-spark-rapids-latest
- emr-7.8.0-spark-rapids-20250228
- emr-7.8.0-java11-latest
- emr-7.8.0-java11-20250228
- emr-7.8.0-java8-latest
- emr-7.8.0-java8-20250228
- emr-7.8.0-spark-rapids-java8-latest
- emr-7.8.0-spark-rapids-java8-20250228
- notebook-spark/emr-7.8.0-latest
- notebook-spark/emr-7.8.0-20250228
- notebook-spark/emr-7.8.0-spark-rapids-latest
- notebook-spark/emr-7.8.0-spark-rapids-20250228
- notebook-spark/emr-7.8.0-java11-latest
- notebook-spark/emr-7.8.0-java11-20250228
- notebook-spark/emr-7.8.0-java8-latest
- notebook-spark/emr-7.8.0-java8-20250228
- notebook-spark/emr-7.8.0-spark-rapids-java8-latest
- notebook-spark/emr-7.8.0-spark-rapids-java8-20250228
- notebook-python/emr-7.8.0-latest
- notebook-python/emr-7.8.0-20250228
- notebook-python/emr-7.8.0-spark-rapids-latest

- notebook-python/emr-7.8.0-spark-rapids-20250228
- notebook-python/emr-7.8.0-java11-latest
- notebook-python/emr-7.8.0-java11-20250228
- notebook-python/emr-7.8.0-java8-latest
- notebook-python/emr-7.8.0-java8-20250228
- notebook-python/emr-7.8.0-spark-rapids-java8-latest
- notebook-python/emr-7.8.0-spark-rapids-java8-20250228
- livy/emr-7.8.0-latest
- livy/emr-7.8.0-20250228
- livy/emr-7.8.0-java11-latest
- livy/emr-7.8.0-java11-20250228
- livy/emr-7.8.0-java8-latest
- livy/emr-7.8.0-java8-20250228

Release notes

Release notes for Amazon EMR on EKS 7.8.0

- **Supported applications** - AWS SDK for Java 2.29.52 and 1.12.780, Apache Spark 3.5.4, Apache Hudi 0.15.0-amzn-5, Apache Iceberg 1.7.1-amzn-1, Delta 3.3.0-amzn-0, Apache Spark RAPIDS 24.12.0-amzn-0, Jupyter Enterprise Gateway 2.6.0, Apache Flink 1.20.0-amzn-2, Flink Operator 1.10.0-amzn-2
- **Supported components** - emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the core-site.xml Hadoop file.
emrfs-site	Change EMRFS settings.

Classifications	Descriptions
spark-metrics	Change values in the <code>metrics.properties</code> Spark file.
spark-defaults	Change values in the <code>spark-defaults.conf</code> Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the <code>hive-site.xml</code> Spark file.
spark-log4j2	Change values in the <code>log4j2.properties</code> Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Changes

The following changes are included with the 7.8.0 release of Amazon EMR on EKS:

- Native-FGAC features, including:
 - Iceberg support to run jobs that perform actions on Non-Lake Formation Tables in a fine-grained access control(FGAC) virtual cluster. (There is a fallback to IAM.)
 - S3 table support
- Spark connect

emr-7.8.0-latest

Release notes: `emr-7.8.0-latest` currently points to `emr-7.8.0-20250228`.

Regions: `emr-7.8.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.8.0:latest`

emr-7.8.0-20250228

Release notes: `emr-7.8.0-20250228` was released in February 2025. This is the initial release of Amazon EMR 7.8.0 (Spark).

Regions: `emr-emr-7.8.0-20250228` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.8.0-20250228`

emr-7.8.0-flink-latest

Release notes: `emr-7.8.0-flink-latest` currently points to `emr-7.8.0-flink-20250228`

Regions: `emr-7.8.0-flink-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.8.0-flink:latest`

emr-7.8.0-flink-20250228

Release notes: `7.8.0-flink-20250228` was released in February 2025. This is the initial release of Amazon EMR 7.8.0 (Flink).

Regions: `emr-7.8.0-flink-20250228` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.8.0-flink:20250228`

Amazon EMR on EKS 7.7.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.7.0 release in general, see [Amazon EMR 7.7.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.7 releases

The following Amazon EMR 7.7.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.7.0-XXXX** release to view more details such as the related container image tag.

Flink releases

The following Amazon EMR 7.7.0 releases are available for Amazon EMR on EKS when you run Flink applications.

- [emr-7.7.0-flink-latest](#)
- [emr-7.7.0-flink-20250131](#)

Spark releases

The following Amazon EMR 7.7.0 releases are available for Amazon EMR on EKS when you run Spark applications.

- [emr-7.7.0-latest](#)
- [emr-7.7.0-20250131](#)
- `emr-7.7.0-spark-rapids-latest`
- `emr-7.7.0-spark-rapids-20250131`
- `emr-7.7.0-java11-latest`
- `emr-7.7.0-java11-20250131`
- `emr-7.7.0-java8-latest`
- `emr-7.7.0-java8-20250131`
- `emr-7.7.0-spark-rapids-java8-latest`

- `emr-7.7.0-spark-rapids-java8-20250131`
- `notebook-spark/emr-7.7.0-latest`
- `notebook-spark/emr-7.7.0-20250131`
- `notebook-spark/emr-7.7.0-spark-rapids-latest`
- `notebook-spark/emr-7.7.0-spark-rapids-20250131`
- `notebook-spark/emr-7.7.0-java11-latest`
- `notebook-spark/emr-7.7.0-java11-20250131`
- `notebook-spark/emr-7.7.0-java8-latest`
- `notebook-spark/emr-7.7.0-java8-20250131`
- `notebook-spark/emr-7.7.0-spark-rapids-java8-latest`
- `notebook-spark/emr-7.7.0-spark-rapids-java8-20250131`
- `notebook-python/emr-7.7.0-latest`
- `notebook-python/emr-7.7.0-20250131`
- `notebook-python/emr-7.7.0-spark-rapids-latest`
- `notebook-python/emr-7.7.0-spark-rapids-20250131`
- `notebook-python/emr-7.7.0-java11-latest`
- `notebook-python/emr-7.7.0-java11-20250131`
- `notebook-python/emr-7.7.0-java8-latest`
- `notebook-python/emr-7.7.0-java8-20250131`
- `notebook-python/emr-7.7.0-spark-rapids-java8-latest`
- `notebook-python/emr-7.7.0-spark-rapids-java8-20250131`
- `livy/emr-7.7.0-latest`
- `livy/emr-7.7.0-20250131`
- `livy/emr-7.7.0-java11-latest`
- `livy/emr-7.7.0-java11-20250131`
- `livy/emr-7.7.0-java8-latest`
- `livy/emr-7.7.0-java8-20250131`

Release notes

Release notes for Amazon EMR on EKS 7.7.0

- **Supported applications** - AWS SDK for Java 2.29.25 and 1.12.779, Apache Spark 3.5.3-amzn-0, Apache Hudi 0.15.0-amzn-3, Apache Iceberg 1.6.1-amzn-2, Delta 3.2.1-amzn-1, Apache Spark RAPIDS 24.10.1-amzn-0, Jupyter Enterprise Gateway 2.6.0, Apache Flink 1.20.0-amzn-0, Flink Operator 1.10.0-amzn-0
- **Supported components** - emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the core-site.xml Hadoop file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the metrics.properties Spark file.
spark-defaults	Change values in the spark-defaults.conf Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the hive-site.xml Spark file.
spark-log4j2	Change values in the log4j2.properties Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Changes

The following changes are included with the 7.7.0 release of Amazon EMR on EKS:

- The Iceberg version in use as of EMR 7.7.0 no longer supports Java 8. Additionally, Iceberg is excluded from the following Java 8 images: `emr-7.7.0-java8-latest` and `emr-7.7.0-spark-rapids-java8-latest`.

emr-7.7.0-latest

Release notes: `emr-7.7.0-latest` currently points to `emr-7.7.0-20250131`.

Regions: `emr-7.7.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.7.0:latest`

emr-7.7.0-20250131

Release notes: `emr-7.7.0-20250131` was released in February 2025. This is the initial release of Amazon EMR 7.7.0 (Spark).

Regions: `emr-emr-7.7.0-20250131` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.7.0-20250131`

emr-7.7.0-flink-latest

Release notes: `emr-7.7.0-flink-latest` currently points to `emr-7.7.0-flink-20250131`

Regions: `emr-7.7.0-flink-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.7.0-flink:latest`

emr-7.7.0-flink-20250131

Release notes: `7.7.0-flink-20250131` was released in February 2025. This is the initial release of Amazon EMR 7.7.0 (Flink).

Regions: `emr-7.7.0-flink-20250131` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.7.0-flink:20250131`

Amazon EMR on EKS 7.6.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.6.0 release in general, see [Amazon EMR 7.6.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.6 releases

The following Amazon EMR 7.6.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.6.0-XXXX** release to view more details such as the related container image tag.

Flink releases

The following Amazon EMR 7.6.0 releases are available for Amazon EMR on EKS when you run Flink applications.

- [emr-7.6.0-flink-latest](#)

- [emr-7.6.0-flink-20241213](#)

Spark releases

The following Amazon EMR 7.6.0 releases are available for Amazon EMR on EKS when you run Spark applications.

- [emr-7.6.0-latest](#)
- [emr-7.6.0-20241213](#)
- emr-7.6.0-spark-rapids-latest
- emr-7.6.0-spark-rapids-20241213
- emr-7.6.0-java11-latest
- emr-7.6.0-java11-20241213
- emr-7.6.0-java8-latest
- emr-7.6.0-java8-20241213
- emr-7.6.0-spark-rapids-java8-latest
- emr-7.6.0-spark-rapids-java8-20241213
- notebook-spark/emr-7.6.0-latest
- notebook-spark/emr-7.6.0-20241213
- notebook-spark/emr-7.6.0-spark-rapids-latest
- notebook-spark/emr-7.6.0-spark-rapids-20241213
- notebook-spark/emr-7.6.0-java11-latest
- notebook-spark/emr-7.6.0-java11-20241213
- notebook-spark/emr-7.6.0-java8-latest
- notebook-spark/emr-7.6.0-java8-20241213
- notebook-spark/emr-7.6.0-spark-rapids-java8-latest
- notebook-spark/emr-7.6.0-spark-rapids-java8-20241213
- notebook-python/emr-7.6.0-latest
- notebook-python/emr-7.6.0-20241213
- notebook-python/emr-7.6.0-spark-rapids-latest
- notebook-python/emr-7.6.0-spark-rapids-20241213
- notebook-python/emr-7.6.0-java11-latest

- notebook-python/emr-7.6.0-java11-20241213
- notebook-python/emr-7.6.0-java8-latest
- notebook-python/emr-7.6.0-java8-20241213
- notebook-python/emr-7.6.0-spark-rapids-java8-latest
- notebook-python/emr-7.6.0-spark-rapids-java8-20241213
- livy/emr-7.6.0-latest
- livy/emr-7.6.0-20241213
- livy/emr-7.6.0-java11-latest
- livy/emr-7.6.0-java11-20241213
- livy/emr-7.6.0-java8-latest
- livy/emr-7.6.0-java8-20241213

Release notes

Release notes for Amazon EMR on EKS 7.6.0

- **Supported applications** - AWS SDK for Java 2.29.25 and 1.12.779, Apache Spark 3.5.3-amzn-0, Apache Hudi 0.15.0-amzn-3, Apache Iceberg 1.6.1-amzn-2, Delta 3.2.1-amzn-1, Apache Spark RAPIDS 24.10.1-amzn-0, Jupyter Enterprise Gateway 2.6.0, Apache Flink 1.20.0-amzn-0, Flink Operator 1.10.0-amzn-0
- **Supported components** - aws-sagemaker-spark-sdk, emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the core-site.xml Hadoop file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the metrics.properties Spark file.

Classifications	Descriptions
spark-defaults	Change values in the spark-defaults.conf Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the hive-site.xml Spark file.
spark-log4j2	Change values in the log4j2.properties Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway jupyter_enterprise_gateway_config.py file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as spark-hive-site.xml. For more information, see [Configure Applications](#).

Notable features

The following features are included with the 7.6.0 release of Amazon EMR on EKS:

- **Monitoring Configuration Support for Apache Spark Operator** – Monitoring configuration lets you easily set up log archiving of your Spark application and operator logs to Amazon S3 or to Amazon CloudWatch. You can choose one or both. Doing so adds a log agent sidecar to your

Spark operator pod, driver, and executor pods, and subsequently forwards these components' logs to your configured sinks. For more information, see [Using monitoring configuration to monitor the Spark Kubernetes operator and Spark jobs](#).

Changes

The following changes are included with the 7.6.0 release of Amazon EMR on EKS:

- No changes for the release.

emr-7.6.0-latest

Release notes: `emr-7.6.0-latest` currently points to `emr-7.6.0-20241213`.

Regions: `emr-7.6.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.6.0:latest`

emr-7.6.0-20241213

Release notes: `7.6.0-20241213` was released in January, 2024. This is the initial release of Amazon EMR 7.6.0 (Spark).

Regions: `emr-7.6.0-20241213` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.6.0:20241213`

emr-7.6.0-flink-latest

Release notes: `emr-7.6.0-flink-latest` currently points to `emr-7.6.0-flink-20241213`

Regions: `emr-7.6.0-flink-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.6.0-flink:latest`

emr-7.6.0-flink-20241213

Release notes: 7.6.0-flink-20241213 was released in January 2024. This is the initial release of Amazon EMR 7.6.0 (Flink).

Regions: `emr-7.6.0-flink-20241213` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.6.0-flink:20241213`

Amazon EMR on EKS 7.5.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.5.0 release in general, see [Amazon EMR 7.5.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.5 releases

The following Amazon EMR 7.5.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.5.0-XXXX** release to view more details such as the related container image tag.

Release notes

Release notes for Amazon EMR on EKS 7.5.0

- **Supported applications** - AWS SDK for Java 2.28.8 and 1.12.772, Apache Spark 3.5.2-amzn-1, Apache Hudi 0.15.0-amzn-1, Apache Iceberg 1.6.1-amzn-0, Delta 3.2.0-amzn-1, Apache Spark RAPIDS 24.08.1-amzn-1, Jupyter Enterprise Gateway 2.6.0, Apache Flink 1.19.1-amzn-1, Flink Operator 1.9.0-amzn-0
- **Supported components** - `aws-sagemaker-spark-sdk`, `emr-ddb`, `emr-goodies`, `emr-s3-select`, `emrfs`, `hadoop-client`, `hudi`, `hudi-spark`, `iceberg`, `spark-kubernetes`.

emr-7.5.0 common vulnerabilities and exposures

The following table lists all CVEs that don't impact EMR on EKS job runs that run on recommended configurations of `emr-7.5.0`. Amazon EMR is dependent on upstream open-source for availability

of fixes and will provide the latest stable version as part of the Core Engine components within 90 days of Amazon EMR verifying the fixes.

This table was updated on August 21, 2026.

CVE ID	Severity	CVE details URL
CVE-2023-39410	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2023-39410
CVE-2024-47561	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2024-47561
CVE-2025-46762	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2025-46762

Amazon EMR on EKS 7.4.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.4.0 release in general, see [Amazon EMR 7.4.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.4 releases

The following Amazon EMR 7.4.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.4.0-XXXX** release to view more details such as the related container image tag.

Release notes

Release notes for Amazon EMR on EKS 7.4.0

- **Supported applications** - AWS SDK for Java 2.25.70 and 1.12.772, Apache Spark 3.5.2-amzn-0, Apache Hudi 0.15.0-amzn-1, Apache Iceberg 1.6.1-amzn-0, Delta 3.2.0-amzn-1, Apache Spark RAPIDS 24.08.1-amzn-0, Jupyter Enterprise Gateway 2.6.0, Apache Flink 1.19.1-amzn-0, Flink Operator 1.9.0-amzn-1
- **Supported components** - aws-sagemaker-spark-sdk, emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.

Amazon EMR on EKS 7.3.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.3.0 release in general, see [Amazon EMR 7.3.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.3 releases

The following Amazon EMR 7.3.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.3.0-XXXX** release to view more details such as the related container image tag.

Flink releases

The following Amazon EMR 7.3.0 releases are available for Amazon EMR on EKS when you run Flink applications.

- [emr-7.3.0-flink-latest](#)
- [emr-7.3.0-flink-29240920](#)

Spark releases

The following Amazon EMR 7.3.0 releases are available for Amazon EMR on EKS when you run Spark applications.

- [emr-7.3.0-latest](#)
- [emr-7.3.0-20240920](#)
- emr-7.3.0-spark-rapids-latest
- emr-7.3.0-spark-rapids-29240920
- emr-7.3.0-java11-latest
- emr-7.3.0-java11-29240920
- emr-7.3.0-java8-latest
- emr-7.3.0-java8-29240920
- emr-7.3.0-spark-rapids-java8-latest
- emr-7.3.0-spark-rapids-java8-29240920
- notebook-spark/emr-7.3.0-latest

- notebook-spark/emr-7.3.0-20240920
- notebook-spark/emr-7.3.0-spark-rapids-latest
- notebook-spark/emr-7.3.0-spark-rapids-29240920
- notebook-spark/emr-7.3.0-java11-latest
- notebook-spark/emr-7.3.0-java11-29240920
- notebook-spark/emr-7.3.0-java8-latest
- notebook-spark/emr-7.3.0-java8-29240920
- notebook-spark/emr-7.3.0-spark-rapids-java8-latest
- notebook-spark/emr-7.3.0-spark-rapids-java8-29240920
- notebook-python/emr-7.3.0-latest
- notebook-python/emr-7.3.0-20240920
- notebook-python/emr-7.3.0-spark-rapids-latest
- notebook-python/emr-7.3.0-spark-rapids-29240920
- notebook-python/emr-7.3.0-java11-latest
- notebook-python/emr-7.3.0-java11-29240920
- notebook-python/emr-7.3.0-java8-latest
- notebook-python/emr-7.3.0-java8-29240920
- notebook-python/emr-7.3.0-spark-rapids-java8-latest
- notebook-python/emr-7.3.0-spark-rapids-java8-29240920
- livy/emr-7.3.0-latest
- livy/emr-7.3.0-20240920
- livy/emr-7.3.0-java11-latest
- livy/emr-7.3.0-java11-29240920
- livy/emr-7.3.0-java8-latest
- livy/emr-7.3.0-java8-29240920

Release notes

Release notes for Amazon EMR on EKS 7.3.0

- **Supported applications** - AWS SDK for Java 2.25.70 and 1.12.747, Apache Spark 3.5.1-amzn-1, Apache Hudi 0.15.0-amzn-0, Apache Iceberg 1.5.2-amzn-0, Delta 3.2.0-amzn-0, Apache Spark

RAPIDS 24.06.1-amzn-0, Jupyter Enterprise Gateway 2.6.0, Apache Flink 1.18.1-amzn-2, Flink Operator 1.9.0-amzn-0

- **Supported components** - aws-sagemaker-spark-sdk, emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the core-site.xml Hadoop file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the metrics.properties Spark file.
spark-defaults	Change values in the spark-defaults.conf Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the hive-site.xml Spark file.
spark-log4j2	Change values in the log4j2.properties Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Notable features

The following features are included with the 7.3.0 release of Amazon EMR on EKS.

- **Application upgrades** – Amazon EMR on EKS now includes [Flink Operator](#) 1.9.0. In addition to other features, the Flink Kubernetes now lets you set CPU and memory quotas for the autoscaler.
- **Apache Iceberg support for Apache Flink** – Apache Iceberg is an open-source high-performance format huge analytic tables. Starting with Amazon EMR 7.3.0, you can use Apache Iceberg tables when you run Apache Flink on Amazon EMR on EKS. For more information, see the Amazon EMR on EKS [Using Apache Iceberg with Amazon EMR on EKS](#).
- **Delta Lake support for Apache Flink** – Delta Lake is a storage layer framework for lakehouse architectures commonly built on Amazon S3. With Amazon EMR 7.3.0 and higher, you can use Delta tables when you run Apache Flink on Amazon EMR on EKS. For more information, see [Using Delta Lake with Amazon EMR on EKS](#).

Changes

The following changes are included with the 7.3.0 release of Amazon EMR on EKS.

- With Amazon EMR on EKS 7.3.0 and higher, Apache Flink now uses Java 17 runtime by default.

emr-7.3.0-latest

Release notes: `emr-7.3.0-latest` currently points to `emr-7.3.0-20240920`.

Regions: `emr-7.3.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.3.0:latest`

emr-7.3.0-20240920

Release notes: `7.3.0-20240920` was released in December, 2023. This is the initial release of Amazon EMR 7.3.0 (Spark).

Regions: `emr-7.3.0-20240920` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.3.0:29240920`

emr-7.3.0 common vulnerabilities and exposures

The following table lists all CVEs that don't impact EMR on EKS job runs that run on recommended configurations of `emr-7.3.0`. Amazon EMR is dependent on upstream open-source for availability of fixes and will provide the latest stable version as part of the Core Engine components within 90 days of Amazon EMR verifying the fixes.

This table was updated on August 21, 2026.

CVE ID	Severity	CVE details URL
CVE-2020-13949	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2020-13949
CVE-2020-28052	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2020-28052
CVE-2021-0341	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2021-0341
CVE-2022-34917	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2022-34917

CVE ID	Severity	CVE details URL
CVE-2022-3509	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2022-3509
CVE-2022-46337	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2022-46337
CVE-2023-25194	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-25194
CVE-2023-2976	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-2976
CVE-2023-3635	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-3635
CVE-2023-52428	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-52428
CVE-2024-25638	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-25638
CVE-2024-31141	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-31141
CVE-2024-36114	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-36114
CVE-2024-51504	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-51504
CVE-2024-7254	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-7254
CVE-2025-24970	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-24970
CVE-2025-27817	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-27817

CVE ID	Severity	CVE details URL
CVE-2025-27818	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2025-27818
CVE-2011-4969	LOW	https://nvd.nist.gov/vuln/detail/CVE-2011-4969
CVE-2018-14040	LOW	https://nvd.nist.gov/vuln/detail/CVE-2018-14040
CVE-2019-8331	LOW	https://nvd.nist.gov/vuln/detail/CVE-2019-8331
CVE-2021-22569	LOW	https://nvd.nist.gov/vuln/detail/CVE-2021-22569
CVE-2021-38153	LOW	https://nvd.nist.gov/vuln/detail/CVE-2021-38153
CVE-2023-42503	LOW	https://nvd.nist.gov/vuln/detail/CVE-2023-42503
CVE-2024-23944	LOW	https://nvd.nist.gov/vuln/detail/CVE-2024-23944
CVE-2024-25710	LOW	https://nvd.nist.gov/vuln/detail/CVE-2024-25710
CVE-2024-26308	LOW	https://nvd.nist.gov/vuln/detail/CVE-2024-26308
CVE-2024-56128	LOW	https://nvd.nist.gov/vuln/detail/CVE-2024-56128
CVE-2025-49128	LOW	https://nvd.nist.gov/vuln/detail/CVE-2025-49128

emr-7.3.0-flink-latest

Release notes: `emr-7.3.0-flink-latest` currently points to `emr-7.3.0-flink-29240920`.

Regions: `emr-7.3.0-flink-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.3.0-flink:latest`

emr-7.3.0-flink-29240920

Release notes: `7.3.0-flink-29240920` was released in December 2023. This is the initial release of Amazon EMR 7.3.0 (Flink).

Regions: `emr-7.3.0-flink-29240920` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.3.0-flink:29240920`

Amazon EMR on EKS 7.2.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.2.0 release in general, see [Amazon EMR 7.2.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.2 releases

The following Amazon EMR 7.2.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.2.0-XXXX** release to view more details such as the related container image tag.

Flink releases

The following Amazon EMR 7.2.0 releases are available for Amazon EMR on EKS when you run Flink applications.

- [emr-7.2.0-flink-latest](#)
- [emr-7.2.0-flink-20240610](#)

Spark releases

The following Amazon EMR 7.2.0 releases are available for Amazon EMR on EKS when you run Spark applications.

- [emr-7.2.0-latest](#)
- [emr-7.2.0-20240610](#)
- emr-7.2.0-spark-rapids-latest
- emr-7.2.0-spark-rapids-20240610
- emr-7.2.0-java11-latest
- emr-7.2.0-java11-20240610
- emr-7.2.0-java8-latest
- emr-7.2.0-java8-20240610
- emr-7.2.0-spark-rapids-java8-latest
- emr-7.2.0-spark-rapids-java8-20240610
- notebook-spark/emr-7.2.0-latest
- notebook-spark/emr-7.2.0-20240610
- notebook-spark/emr-7.2.0-spark-rapids-latest
- notebook-spark/emr-7.2.0-spark-rapids-20240610
- notebook-spark/emr-7.2.0-java11-latest
- notebook-spark/emr-7.2.0-java11-20240610
- notebook-spark/emr-7.2.0-java8-latest
- notebook-spark/emr-7.2.0-java8-20240610
- notebook-spark/emr-7.2.0-spark-rapids-java8-latest
- notebook-spark/emr-7.2.0-spark-rapids-java8-20240610
- notebook-python/emr-7.2.0-latest
- notebook-python/emr-7.2.0-20240610
- notebook-python/emr-7.2.0-spark-rapids-latest
- notebook-python/emr-7.2.0-spark-rapids-20240610
- notebook-python/emr-7.2.0-java11-latest
- notebook-python/emr-7.2.0-java11-20240610

- `notebook-python/emr-7.2.0-java8-latest`
- `notebook-python/emr-7.2.0-java8-20240610`
- `notebook-python/emr-7.2.0-spark-rapids-java8-latest`
- `notebook-python/emr-7.2.0-spark-rapids-java8-20240610`
- `livy/emr-7.2.0-latest`
- `livy/emr-7.2.0-20240610`
- `livy/emr-7.2.0-java11-latest`
- `livy/emr-7.2.0-java11-20240610`
- `livy/emr-7.2.0-java8-latest`
- `livy/emr-7.2.0-java8-20240610`

Release notes

Release notes for Amazon EMR on EKS 7.2.0

- **Supported applications** - AWS SDK for Java 2.23.18 and 1.12.705, Apache Spark 3.5.1-amzn-1, Apache Hudi 0.14.1-amzn-0, Apache Iceberg 1.5.0-amzn-0, Delta 3.1.0, Apache Spark RAPIDS 24.02.0-amzn-1, Jupyter Enterprise Gateway 2.6.0, Apache Flink 1.18.1-amzn-0, Flink Operator 1.8.0-amzn-1
- **Supported components** - `aws-sagemaker-spark-sdk`, `emr-ddb`, `emr-goodies`, `emr-s3-select`, `emrfs`, `hadoop-client`, `hudi`, `hudi-spark`, `iceberg`, `spark-kubernetes`.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
<code>core-site</code>	Change values in the <code>core-site.xml</code> Hadoop file.
<code>emrfs-site</code>	Change EMRFS settings.
<code>spark-metrics</code>	Change values in the <code>metrics.properties</code> Spark file.

Classifications	Descriptions
spark-defaults	Change values in the spark-defaults.conf Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the hive-site.xml Spark file.
spark-log4j2	Change values in the log4j2.properties Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway jupyter_enterprise_gateway_config.py file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as spark-hive-site.xml. For more information, see [Configure Applications](#).

Notable features

The following features are included with the 7.2.0 release of Amazon EMR on EKS.

- **Application upgrades** – Amazon EMR on EKS 7.2.0 application upgrades include Spark 3.5.1, Flink 1.18.1, and [Flink Operator](#) 1.8.0.

- [Autoscaler for Flink updates](#) – The 7.2.0 release uses the open source configuration `job.autoscaler.restart.time-tracking.enabled` to enable rescale time estimation, so you no longer have to manually assign empirical values to restart time. If you run 7.1.0 or lower, you can still use Amazon EMR autoscaling.
- [Apache Hudi integration Apache Flink on Amazon EMR on EKS](#) – This release adds an integration between Apache Hudi and Apache Flink, so you can use the Flink Kubernetes operator to run Hudi jobs. Hudi lets you use record-level operations that you can use to simplify data management and data pipeline development.
- [Amazon S3 Express One Zone integration with Amazon EMR on EKS](#) – With 7.2.0 and higher, you can upload data into the S3 Express One Zone with Amazon EMR on EKS. S3 Express One Zone is a high-performance, single-zone Amazon S3 storage class that delivers consistent, single-digit millisecond data access for most latency-sensitive applications. At the time of its release, S3 Express One Zone delivers the lowest latency and highest performance cloud object storage in Amazon S3.
- [Support for default configurations in the Spark operator](#) – Spark operator on Amazon EKS now supports the same default configurations as the start job run model on Amazon EMR on EKS for 7.2.0 and higher. This means that features such as Amazon S3 and EMRFS no longer require manual configurations in the yaml file.

emr-7.2.0-latest

Release notes: `emr-7.2.0-latest` currently points to `emr-7.2.0-20240610`.

Regions: `emr-7.2.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.2.0:latest`

emr-7.2.0-20240610

Release notes: `7.2.0-20240610` was released in December, 2023. This is the initial release of Amazon EMR 7.2.0 (Spark).

Regions: `emr-7.2.0-20240610` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.2.0:20240610`

emr-7.2.0 common vulnerabilities and exposures

The following table lists all CVEs that don't impact EMR on EKS job runs that run on recommended configurations of emr-7.2.0. Amazon EMR is dependent on upstream open-source for availability of fixes and will provide the latest stable version as part of the Core Engine components within 90 days of Amazon EMR verifying the fixes.

This table was updated on August 21, 2026.

CVE ID	Severity	CVE details URL
CVE-2019-10202	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2019-10202
CVE-2022-46337	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2022-46337
CVE-2019-0205	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2019-0205
CVE-2019-10172	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2019-10172
CVE-2020-13949	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2020-13949
CVE-2020-28052	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2020-28052
CVE-2021-22569	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2021-22569
CVE-2021-22570	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2021-22570
CVE-2021-31684	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2021-31684
CVE-2021-33503	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2021-33503

CVE ID	Severity	CVE details URL
CVE-2022-3509	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2022-3509
CVE-2022-3510	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2022-3510
CVE-2022-40897	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2022-40897
CVE-2022-45688	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2022-45688
CVE-2022-46751	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2022-46751
CVE-2023-0286	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-0286
CVE-2023-1370	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-1370
CVE-2023-39410	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-39410
CVE-2023-43642	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-43642
CVE-2023-5072	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-5072
CVE-2023-50782	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-50782
CVE-2024-21634	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-21634
CVE-2024-25638	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-25638

CVE ID	Severity	CVE details URL
CVE-2024-25710	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-25710
CVE-2024-26308	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-26308
CVE-2024-36114	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-36114
CVE-2024-6345	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-6345
CVE-2024-7254	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2024-7254
CVE-2018-10237	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2018-10237
CVE-2021-29425	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2021-29425
CVE-2022-3171	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2022-3171
CVE-2023-2976	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2023-2976
CVE-2023-33201	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2023-33201
CVE-2023-33202	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2023-33202
CVE-2023-3635	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2023-3635
CVE-2023-42503	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2023-42503

CVE ID	Severity	CVE details URL
CVE-2023-52428	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2023-52428
CVE-2024-23944	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2024-23944
CVE-2024-29025	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2024-29025
CVE-2024-29131	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2024-29131
CVE-2024-29133	MEDIUM	https://nvd.nist.gov/vuln/detail/CVE-2024-29133
CVE-2020-8908	LOW	https://nvd.nist.gov/vuln/detail/CVE-2020-8908

emr-7.2.0-flink-latest

Release notes: `emr-7.2.0-flink-latest` currently points to `emr-7.2.0-flink-20240610`.

Regions: `emr-7.2.0-flink-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.2.0-flink:latest`

emr-7.2.0-flink-20240610

Release notes: `7.2.0-flink-20240610` was released in December 2023. This is the initial release of Amazon EMR 7.2.0 (Flink).

Regions: `emr-7.2.0-flink-20240610` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.2.0-flink:20240610`

Amazon EMR on EKS 7.1.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.1.0 release in general, see [Amazon EMR 7.1.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.1 releases

The following Amazon EMR 7.1.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.1.0-XXXX** release to view more details such as the related container image tag.

Flink releases

The following Amazon EMR 7.1.0 releases are available for Amazon EMR on EKS when you run Flink applications.

- [emr-7.1.0-flink-latest](#)
- [emr-7.1.0-flink-20240321](#)

Spark releases

The following Amazon EMR 7.1.0 releases are available for Amazon EMR on EKS when you run Spark applications.

- [emr-7.1.0-latest](#)
- [emr-7.1.0-20240321](#)
- emr-7.1.0-spark-rapids-latest
- emr-7.1.0-spark-rapids-20240321
- emr-7.1.0-java11-latest
- emr-7.1.0-java11-20240321
- emr-7.1.0-java8-latest
- emr-7.1.0-java8-20240321
- emr-7.1.0-spark-rapids-java8-latest
- emr-7.1.0-spark-rapids-java8-20240321
- notebook-spark/emr-7.1.0-latest

- notebook-spark/emr-7.1.0-20240321
- notebook-spark/emr-7.1.0-spark-rapids-latest
- notebook-spark/emr-7.1.0-spark-rapids-20240321
- notebook-spark/emr-7.1.0-java11-latest
- notebook-spark/emr-7.1.0-java11-20240321
- notebook-spark/emr-7.1.0-java8-latest
- notebook-spark/emr-7.1.0-java8-20240321
- notebook-spark/emr-7.1.0-spark-rapids-java8-latest
- notebook-spark/emr-7.1.0-spark-rapids-java8-20240321
- notebook-python/emr-7.1.0-latest
- notebook-python/emr-7.1.0-20240321
- notebook-python/emr-7.1.0-spark-rapids-latest
- notebook-python/emr-7.1.0-spark-rapids-20240321
- notebook-python/emr-7.1.0-java11-latest
- notebook-python/emr-7.1.0-java11-20240321
- notebook-python/emr-7.1.0-java8-latest
- notebook-python/emr-7.1.0-java8-20240321
- notebook-python/emr-7.1.0-spark-rapids-java8-latest
- notebook-python/emr-7.1.0-spark-rapids-java8-20240321
- livy/emr-7.1.0-latest
- livy/emr-7.1.0-20240321
- livy/emr-7.1.0-java11-latest
- livy/emr-7.1.0-java11-20240321
- livy/emr-7.1.0-java8-latest
- livy/emr-7.1.0-java8-20240321

Release notes

Release notes for Amazon EMR on EKS 7.1.0

- **Supported applications** - AWS SDK for Java 2.23.18 and 1.12.656, Apache Spark 3.5.0-amzn-1, Apache Hudi 0.14.1-amzn-0, Apache Iceberg 1.4.3-amzn-0, Delta 3.0.0, Apache Spark RAPIDS

23.10.0-amzn-1, Jupyter Enterprise Gateway 2.6.0, Apache Flink 1.18.1-amzn-0, Flink Operator 1.6.1-amzn-1

- **Supported components** - aws-sagemaker-spark-sdk, emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the core-site.xml Hadoop file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the metrics.properties Spark file.
spark-defaults	Change values in the spark-defaults.conf Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the hive-site.xml Spark file.
spark-log4j2	Change values in the log4j2.properties Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Notable features

The following features are included with the 7.1.0 release of Amazon EMR on EKS.

- [Apache Livy support for Amazon EMR on EKS](#) – With Amazon EMR on EKS releases 7.1.0 and higher, you can use Apache Livy on an Amazon EKS cluster to create an Apache Livy REST interface to submit Spark jobs or snippets of Spark code. Doing so lets you retrieve results synchronously and asynchronously, while still leveraging Amazon EMR on EKS benefits, such as Amazon EMR-optimized Spark runtime, SSL-enabled Livy endpoints, and a programmatic set-up experience.

emr-7.1.0-latest

Release notes: `emr-7.1.0-latest` currently points to `emr-7.1.0-20240321`.

Regions: `emr-7.1.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.1.0:latest`

emr-7.1.0-20240321

Release notes: 7.1.0-20240321 was released in December, 2023. This is the initial release of Amazon EMR 7.1.0 (Spark).

Regions: `emr-7.1.0-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.1.0:20240321`

emr-7.1.0 common vulnerabilities and exposures

The following table lists all CVEs that don't impact EMR on EKS job runs that run on recommended configurations of `emr-7.1.0`. Amazon EMR is dependent on upstream open-source for availability of fixes and will provide the latest stable version as part of the Core Engine components within 90 days of Amazon EMR verifying the fixes.

This table was updated on August 21, 2026.

CVE ID	Severity	CVE details URL
CVE-2021-41186	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2021-41186

emr-7.1.0-flink-latest

Release notes: `emr-7.1.0-flink-latest` currently points to `emr-7.1.0-flink-20240321`.

Regions: `emr-7.1.0-flink-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.1.0-flink:latest`

emr-7.1.0-flink-20240321

Release notes: 7.1.0-flink-20240321 was released in December 2023. This is the initial release of Amazon EMR 7.1.0 (Flink).

Regions: `emr-7.1.0-flink-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.1.0-flink:20240321`

Amazon EMR on EKS 7.0.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 7.0.0 release in general, see [Amazon EMR 7.0.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 7.0 releases

The following Amazon EMR 7.0.0 releases are available for Amazon EMR on EKS. Select a specific **emr-7.0.0-XXXX** release to view more details such as the related container image tag.

Flink releases

The following Amazon EMR 7.0.0 releases are available for Amazon EMR on EKS when you run Flink applications.

- [emr-7.0.0-flink-latest](#)
- [emr-7.0.0-flink-2024321](#)
- [emr-7.0.0-flink-20231211](#)

Spark releases

The following Amazon EMR 7.0.0 releases are available for Amazon EMR on EKS when you run Spark applications.

- [emr-7.0.0-latest](#)
- [emr-7.0.0-20231211](#)
- `emr-7.0.0-spark-rapids-latest`
- `emr-7.0.0-spark-rapids-20231211`
- `emr-7.0.0-java11-latest`
- `emr-7.0.0-java11-20231211`
- `emr-7.0.0-java8-latest`
- `emr-7.0.0-java8-20231211`
- `emr-7.0.0-spark-rapids-java8-latest`

- `emr-7.0.0-spark-rapids-java8-20231211`
- `notebook-spark/emr-7.0.0-latest`
- `notebook-spark/emr-7.0.0-20231211`
- `notebook-spark/emr-7.0.0-spark-rapids-latest`
- `notebook-spark/emr-7.0.0-spark-rapids-20231211`
- `notebook-spark/emr-7.0.0-java11-latest`
- `notebook-spark/emr-7.0.0-java11-20231211`
- `notebook-spark/emr-7.0.0-java8-latest`
- `notebook-spark/emr-7.0.0-java8-20231211`
- `notebook-spark/emr-7.0.0-spark-rapids-java8-latest`
- `notebook-spark/emr-7.0.0-spark-rapids-java8-20231211`
- `notebook-python/emr-7.0.0-latest`
- `notebook-python/emr-7.0.0-20231211`
- `notebook-python/emr-7.0.0-spark-rapids-latest`
- `notebook-python/emr-7.0.0-spark-rapids-20231211`
- `notebook-python/emr-7.0.0-java11-latest`
- `notebook-python/emr-7.0.0-java11-20231211`
- `notebook-python/emr-7.0.0-java8-latest`
- `notebook-python/emr-7.0.0-java8-20231211`
- `notebook-python/emr-7.0.0-spark-rapids-java8-latest`
- `notebook-python/emr-7.0.0-spark-rapids-java8-20231211`

Release notes

Release notes for Amazon EMR on EKS 7.0.0

- **Supported applications** - AWS SDK for Java 2.20.160-amzn-0 and 1.12.595, Apache Spark 3.5.0-amzn-0, Apache Flink 1.18.0-amzn-0, Flink Operator 1.6.1, Apache Hudi 0.14.0-amzn-1, Apache Iceberg 1.4.2-amzn-0, Delta 3.0.0, Apache Spark RAPIDS 23.10.0-amzn-0, Jupyter Enterprise Gateway 2.6.0
- **Supported components** - `aws-sagemaker-spark-sdk`, `emr-ddb`, `emr-goodies`, `emr-s3-select`, `emrfs`, `hadoop-client`, `hudi`, `hudi-spark`, `iceberg`, `spark-kubernetes`.

- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the <code>core-site.xml</code> Hadoop file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the <code>metrics.properties</code> Spark file.
spark-defaults	Change values in the <code>spark-defaults.conf</code> Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the <code>hive-site.xml</code> Spark file.
spark-log4j	Change values in the <code>log4j2.properties</code> Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Notable features

The following features are included with the 7.0 release of Amazon EMR on EKS.

- **Application upgrades** – Amazon EMR on EKS 7.0.0 application upgrades include Spark 3.5, Flink 1.18, and [Flink Operator](#) 1.6.1.
- **Flink Autoscaler parameter auto-tuning** – The default parameters that Flink Autoscaler uses for its scaling calculations might not be the optimal value for a given job. Amazon EMR on EKS 7.0.0 uses historical trends of specific captured metrics to calculate the optimal parameter tailored for the job.

Changes

The following changes are included with the 7.0 release of Amazon EMR on EKS.

- **Amazon Linux 2023** – With Amazon EMR on EKS 7.0.0 and higher, all container images are based on Amazon Linux 2023.
- **Spark uses Java 17 as default runtime** – Amazon EMR on EKS 7.0.0 Spark uses Java 17 as default runtime. If you need to, you can switch to use Java 8 or Java 11 with the corresponding release label as provided in the [Amazon EMR on EKS 7.0 releases](#) list.

emr-7.0.0-latest

Release notes: `emr-7.0.0-latest` currently points to `emr-7.0.0-2024321`.

Regions: `emr-7.0.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.0.0:latest`

emr-7.0.0-2024321

Release notes: 7.0.0-2024321 was released on March 11, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-7.0.0-2024321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.0.0:2024321`

emr-7.0.0-20231211

Release notes: 7.0.0-20231211 was released in December, 2023. This is the initial release of Amazon EMR 7.0.0 (Spark).

Regions: `emr-7.0.0-20231211` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.0.0:20231211`

emr-7.0.0 common vulnerabilities and exposures

The following table lists all CVEs that don't impact EMR on EKS job runs that run on recommended configurations of `emr-7.0.0`. Amazon EMR is dependent on upstream open-source for availability of fixes and will provide the latest stable version as part of the Core Engine components within 90 days of Amazon EMR verifying the fixes.

This table was updated on August 21, 2026.

CVE ID	Severity	CVE details URL
CVE-2019-10202	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2019-10202
CVE-2022-46337	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2022-46337
CVE-2023-44981	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2023-44981

CVE ID	Severity	CVE details URL
CVE-2024-32888	CRITICAL	https://nvd.nist.gov/vuln/detail/CVE-2024-32888
CVE-2013-2192	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2013-2192
CVE-2015-5237	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2015-5237
CVE-2015-7430	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2015-7430
CVE-2015-7431	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2015-7431
CVE-2016-5001	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2016-5001
CVE-2017-15288	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2017-15288
CVE-2017-3161	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2017-3161
CVE-2017-3162	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2017-3162
CVE-2018-21234	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2018-21234
CVE-2019-17543	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2019-17543
CVE-2019-20478	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2019-20478
CVE-2020-13949	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2020-13949

CVE ID	Severity	CVE details URL
CVE-2020-35864	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2020-35864
CVE-2020-8908	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2020-8908
CVE-2021-22570	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2021-22570
CVE-2022-32511	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2022-32511
CVE-2023-33546	HIGH	https://nvd.nist.gov/vuln/detail/CVE-2023-33546

emr-7.0.0-flink-latest

Release notes: `emr-7.0.0-flink-latest` currently points to `emr-7.0.0-flink-2024321`.

Regions: `emr-7.0.0-flink-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.0.0-flink:latest`

emr-7.0.0-flink-2024321

Release notes: `7.0.0-flink-2024321` was released on March 11, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-7.0.0-flink-2024321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.0.0-flink:2024321`

emr-7.0.0-flink-20231211

Release notes: 7.0.0-flink-20231211 was released in December 2023. This is the initial release of Amazon EMR 7.0.0 (Flink).

Regions: `emr-7.0.0-flink-20231211` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-7.0.0-flink:20231211`

Amazon EMR on EKS 6.15.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 6.15.0 release in general, see [Amazon EMR 6.15.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 6.15 releases

The following Amazon EMR 6.15.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.15.0-XXXX** release to view more details such as the related container image tag.

Flink releases

The following Amazon EMR 6.15.0 releases are available for Amazon EMR on EKS when you run Flink applications.

- [emr-6.15.0-flink-latest](#)
- [emr-6.15.0-flink-20240105](#)
- [emr-6.15.0-flink-20231109](#)

Spark releases

The following Amazon EMR 6.15.0 releases are available for Amazon EMR on EKS when you run Spark applications.

- [emr-6.15.0-latest](#)
- [emr-6.15.0-20231109](#)

- `emr-6.15.0-spark-rapids-latest`
- `emr-6.15.0-spark-rapids-20231109`
- `emr-6.15.0-java11-latest`
- `emr-6.15.0-java11-20231109`
- `emr-6.15.0-java17-latest`
- `emr-6.15.0-java17-20231109`
- `emr-6.15.0-java17-al2023-latest`
- `emr-6.15.0-java17-al2023-20231109`
- `emr-6.15.0-spark-rapids-java17-latest`
- `emr-6.15.0-spark-rapids-java17-20231109`
- `emr-6.15.0-spark-rapids-java17-al2023-latest`
- `emr-6.15.0-spark-rapids-java17-al2023-20231109`
- `notebook-spark/emr-6.15.0-latest`
- `notebook-spark/emr-6.15.0-20231109`
- `notebook-spark/emr-6.15.0-spark-rapids-latest`
- `notebook-spark/emr-6.15.0-spark-rapids-20231109`
- `notebook-spark/emr-6.15.0-java11-latest`
- `notebook-spark/emr-6.15.0-java11-20231109`
- `notebook-spark/emr-6.15.0-java17-latest`
- `notebook-spark/emr-6.15.0-java17-20231109`
- `notebook-spark/emr-6.15.0-java17-al2023-latest`
- `notebook-spark/emr-6.15.0-java17-al2023-20231109`
- `notebook-python/emr-6.15.0-latest`
- `notebook-python/emr-6.15.0-20231109`
- `notebook-python/emr-6.15.0-spark-rapids-latest`
- `notebook-python/emr-6.15.0-spark-rapids-20231109`
- `notebook-python/emr-6.15.0-java11-latest`
- `notebook-python/emr-6.15.0-java11-20231109`
- `notebook-python/emr-6.15.0-java17-latest`
- `notebook-python/emr-6.15.0-java17-20231109`

- notebook-python/emr-6.15.0-java17-al2023-latest
- notebook-python/emr-6.15.0-java17-al2023-20231109

Release notes

Release notes for Amazon EMR on EKS 6.15.0

- **Supported applications** - AWS SDK for Java 1.12.569, Apache Spark 3.4.1-amzn-2, Apache Flink 1.17.1-amzn-1, Apache Hudi 0.14.0-amzn-0, Apache Iceberg 1.4.0-amzn-0, Delta 2.4.0, Apache Spark RAPIDS 23.08.01-amzn-0, Jupyter Enterprise Gateway 2.6.0
- **Supported components** - aws-sagemaker-spark-sdk, emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the core-site.xml Hadoop file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the metrics.properties Spark file.
spark-defaults	Change values in the spark-defaults.conf Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the hive-site.xml Spark file.
spark-log4j	Change values in the log4j2.properties Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Notable features

The following features are included with the 6.15 release of Amazon EMR on EKS.

- [Amazon EMR on EKS with Apache Flink](#) - With Amazon EMR on EKS 6.15.0, you can run your Apache Flink-based application along with other types of applications on the same Amazon EKS cluster. This helps improve resource utilization and simplify infrastructure management. You can leverage Spot Instances in a Flink application with graceful decommission, and achieve faster restart times with fine-grained recovery and task-local recovery with Amazon EBS. Accessibility and monitoring features include the ability to launch a Flink application with jars that are stored in Amazon S3, access to the AWS Glue Data Catalog, monitoring integration with Amazon S3 and Amazon CloudWatch, and container log rotation.

emr-6.15.0-latest

Release notes: `emr-6.15.0-latest` currently points to `emr-6.15.0-20240105`.

Regions: `emr-6.15.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.15.0:latest`

emr-6.15.0-20240105

Release notes: 6.15.0-20240105 was released on January 17, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.15.0-20240105` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.15.0:20240105`

emr-6.15.0-20231109

Release notes: 6.15.0-20231109 was released on November 17, 2023. This is the initial release of Amazon EMR 6.15.0.

Regions: `emr-6.15.0-20231109` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.15.0:20231109`

emr-6.15.0-flink-latest

Release notes: `emr-6.15.0-flink-latest` currently points to `emr-6.15.0-flink-20240105`.

Regions: `emr-6.15.0-flink-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.15.0-flink:latest`

emr-6.15.0-flink-20240105

Release notes: 6.15.0-flink-20240105 was released on January 17, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.15.0-flink-20240105` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.15.0-flink:20240105`

emr-6.15.0-flink-20231109

Release notes: 6.15.0-flink-20231109 was released on November 17, 2023. This is the initial release of Amazon EMR 6.15.0.

Regions: emr-6.15.0-flink-20231109 is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: emr-6.15.0-flink:20231109

Amazon EMR on EKS 6.14.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 6.14.0 release in general, see [Amazon EMR 6.14.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 6.14 releases

The following Amazon EMR 6.14.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.14.0-XXXX** release to view more details such as the related container image tag.

- [emr-6.14.0-latest](#)
- [emr-6.14.0-20231005](#)
- emr-6.14.0-spark-rapids-latest
- emr-6.14.0-spark-rapids-20231005
- emr-6.14.0-java11-latest
- emr-6.14.0-java11-20231005
- emr-6.14.0-java17-latest
- emr-6.14.0-java17-20231005
- emr-6.14.0-java17-al2023-latest
- emr-6.14.0-java17-al2023-20231005
- emr-6.14.0-spark-rapids-java17-latest
- emr-6.14.0-spark-rapids-java17-20231005
- emr-6.14.0-spark-rapids-java17-al2023-latest
- emr-6.14.0-spark-rapids-java17-al2023-20231005

- notebook-spark/emr-6.14.0-latest
- notebook-spark/emr-6.14.0-20231005
- notebook-spark/emr-6.14.0-spark-rapids-latest
- notebook-spark/emr-6.14.0-spark-rapids-20231005
- notebook-spark/emr-6.14.0-java11-latest
- notebook-spark/emr-6.14.0-java11-20231005
- notebook-spark/emr-6.14.0-java17-latest
- notebook-spark/emr-6.14.0-java17-20231005
- notebook-spark/emr-6.14.0-java17-al2023-latest
- notebook-spark/emr-6.14.0-java17-al2023-20231005
- notebook-python/emr-6.14.0-latest
- notebook-python/emr-6.14.0-20231005
- notebook-python/emr-6.14.0-spark-rapids-latest
- notebook-python/emr-6.14.0-spark-rapids-20231005
- notebook-python/emr-6.14.0-java11-latest
- notebook-python/emr-6.14.0-java11-20231005
- notebook-python/emr-6.14.0-java17-latest
- notebook-python/emr-6.14.0-java17-20231005
- notebook-python/emr-6.14.0-java17-al2023-latest
- notebook-python/emr-6.14.0-java17-al2023-20231005

Release notes

Release notes for Amazon EMR on EKS 6.14.0

- **Supported applications** - AWS SDK for Java 1.12.543, Apache Spark 3.4.1-amzn-1, Apache Hudi 0.13.1-amzn-2, Apache Iceberg 1.3.0-amzn-0, Delta 2.4.0, Apache Spark RAPIDS 23.06.0-amzn-2, Jupyter Enterprise Gateway 2.7.0
- **Supported components** - aws-sagemaker-spark-sdk, emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the <code>core-site.xml</code> Hadoop file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the <code>metrics.properties</code> Spark file.
spark-defaults	Change values in the <code>spark-defaults.conf</code> Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the <code>hive-site.xml</code> Spark file.
spark-log4j	Change values in the <code>log4j2.properties</code> Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Notable features

The following features are included with the 6.14 release of Amazon EMR on EKS.

- **[Apache Livy](#) support** - Amazon EMR on EKS now supports Apache Livy with `spark-submit`.

emr-6.14.0-latest

Release notes: `emr-6.14.0-latest` currently points to `emr-6.14.0-20231005`.

Regions: `emr-6.14.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.14.0:latest`

emr-6.14.0-20231005

Release notes: `6.14.0-20231005` was released on October 17, 2023. This is the initial release of Amazon EMR 6.14.0.

Regions: `emr-6.14.0-20231005` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.14.0:20231005`

Amazon EMR on EKS 6.13.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 6.13.0 release in general, see [Amazon EMR 6.13.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 6.13 releases

The following Amazon EMR 6.13.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.13.0-XXXX** release to view more details such as the related container image tag.

- [emr-6.13.0-latest](#)
- [emr-6.13.0-20230814](#)
- emr-6.13.0-spark-rapids-latest
- emr-6.13.0-spark-rapids-20230814
- emr-6.13.0-java11-latest
- emr-6.13.0-java11-20230814
- emr-6.13.0-java17-latest
- emr-6.13.0-java17-20230814
- emr-6.13.0-java17-al2023-latest
- emr-6.13.0-java17-al2023-20230814
- emr-6.13.0-spark-rapids-java17-latest
- emr-6.13.0-spark-rapids-java17-20230814
- emr-6.13.0-spark-rapids-java17-al2023-latest
- emr-6.13.0-spark-rapids-java17-al2023-20230814
- notebook-spark/emr-6.13.0-latest
- notebook-spark/emr-6.13.0-20230814
- notebook-spark/emr-6.13.0-spark-rapids-latest
- notebook-spark/emr-6.13.0-spark-rapids-20230814
- notebook-spark/emr-6.13.0-java11-latest
- notebook-spark/emr-6.13.0-java11-20230814
- notebook-spark/emr-6.13.0-java17-latest
- notebook-spark/emr-6.13.0-java17-20230814
- notebook-spark/emr-6.13.0-java17-al2023-latest
- notebook-spark/emr-6.13.0-java17-al2023-20230814
- notebook-python/emr-6.13.0-latest
- notebook-python/emr-6.13.0-20230814

- notebook-python/emr-6.13.0-spark-rapids-latest
- notebook-python/emr-6.13.0-spark-rapids-20230814
- notebook-python/emr-6.13.0-java11-latest
- notebook-python/emr-6.13.0-java11-20230814
- notebook-python/emr-6.13.0-java17-latest
- notebook-python/emr-6.13.0-java17-20230814
- notebook-python/emr-6.13.0-java17-al2023-latest
- notebook-python/emr-6.13.0-java17-al2023-20230814

Release notes

Release notes for Amazon EMR on EKS 6.13.0

- **Supported applications** - AWS SDK for Java 1.12.513, Apache Spark 3.4.1-amzn-0, Apache Hudi 0.13.1-amzn-0, Apache Iceberg 1.3.0-amzn-0, Delta 2.4.0, Apache Spark RAPIDS 23.06.0-amzn-1, Jupyter Enterprise Gateway 2.6.0.amzn
- **Supported components** - aws-sagemaker-spark-sdk, emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the core-site.xml Hadoop file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the metrics.properties Spark file.
spark-defaults	Change values in the spark-defaults.conf Spark file.
spark-env	Change values in the Spark environment.

Classifications	Descriptions
spark-hive-site	Change values in the <code>hive-site.xml</code> Spark file.
spark-log4j	Change values in the <code>log4j2.properties</code> Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Notable features

The following features are included with the 6.13 release of Amazon EMR on EKS.

- **Amazon Linux 2023** - With Amazon EMR on EKS 6.13 and higher, you can launch Spark with AL2023 as operating system together with Java 17 runtime. To do this, use release label with `al2023` in its name. For example: `emr-6.13.0-java17-al2023-latest`. We recommend that you validate and run performance tests before you move your production workloads to AL2023 and Java 17.
- [Amazon EMR on EKS with Apache Flink](#) (public preview) - Amazon EMR on EKS releases 6.13 and higher support Apache Flink, available in public preview. With this launch, you can run your

Apache Flink-based application along with other types of applications on the same Amazon EKS cluster. This helps improve resource utilization and simplify infrastructure management. If you already run big data frameworks on Amazon EKS, you can now let Amazon EMR automate your provisioning and management.

emr-6.13.0-latest

Release notes: `emr-6.13.0-latest` currently points to `emr-6.13.0-20230814`.

Regions: `emr-6.13.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.13.0:latest`

emr-6.13.0-20230814

Release notes: `6.13.0-20230814` was released on September 7, 2023. This is the initial release of Amazon EMR 6.13.0.

Regions: `emr-6.13.0-20230814` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.13.0:20230814`

Amazon EMR on EKS 6.12.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 6.12.0 release in general, see [Amazon EMR 6.12.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 6.12 releases

The following Amazon EMR 6.12.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.12.0-XXXX** release to view more details such as the related container image tag.

- [emr-6.12.0-latest](#)
- [emr-6.12.0-20240321](#)

- [emr-6.12.0-20230701](#)
- emr-6.12.0-spark-rapids-latest
- emr-6.12.0-spark-rapids-20230701
- emr-6.12.0-java11-latest
- emr-6.12.0-java11-20230701
- emr-6.12.0-java17-latest
- emr-6.12.0-java17-20230701
- emr-6.12.0-spark-rapids-java17-latest
- emr-6.12.0-spark-rapids-java17-20230701
- notebook-spark/emr-6.12.0-latest
- notebook-spark/emr-6.12.0-20230701
- notebook-spark/emr-6.12.0-spark-rapids-latest
- notebook-spark/emr-6.12.0-spark-rapids-20230701
- notebook-python/emr-6.12.0-latest
- notebook-python/emr-6.12.0-20230701
- notebook-python/emr-6.12.0-spark-rapids-latest
- notebook-python/emr-6.12.0-spark-rapids-20230701

Release notes

Release notes for Amazon EMR on EKS 6.12.0

- **Supported applications** - AWS SDK for Java 1.12.490, Apache Spark 3.4.0-amzn-0, Apache Hudi 0.13.1-amzn-0, Apache Iceberg 1.3.0-amzn-0, Delta 2.4.0, Apache Spark RAPIDS 23.06.0-amzn-0, Jupyter Enterprise Gateway 2.6.0
- **Supported components** - aws-sagemaker-spark-sdk, emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the <code>core-site.xml</code> Hadoop file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the <code>metrics.properties</code> Spark file.
spark-defaults	Change values in the <code>spark-defaults.conf</code> Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the <code>hive-site.xml</code> Spark file.
spark-log4j	Change values in the <code>log4j2.properties</code> Spark file.
emr-job-submitter	Configuration for job submitter pod .

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Notable features

The following features are included with the 6.12 release of Amazon EMR on EKS.

- **Java 17** - With Amazon EMR on EKS 6.12 and higher, you can launch Spark with Java 17 runtime. To do this, pass `emr-6.12.0-java17-latest` as a release label. We recommend that you validate and run performance tests before you move your production workloads from earlier versions of the Java image to the Java 17 image.

emr-6.12.0-latest

Release notes: `emr-6.12.0-latest` currently points to `emr-6.12.0-20240321`.

Regions: `emr-6.12.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.12.0:latest`

emr-6.12.0-20240321

Release notes: `6.12.0-20240321` was released on March 11, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.12.0-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.12.0:20240321`

emr-6.12.0-20230701

Release notes: `6.12.0-20230701` was released on July 1, 2023. This is the initial release of Amazon EMR 6.12.0.

Regions: `emr-6.12.0-20230701` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.12.0:20230701`

Amazon EMR on EKS 6.11.0 releases

This page describes the new and updated functionality for Amazon EMR that is specific to the Amazon EMR on EKS deployment. For details about Amazon EMR running on Amazon EC2 and about the Amazon EMR 6.11.0 release in general, see [Amazon EMR 6.11.0](#) in the *Amazon EMR Release Guide*.

Amazon EMR on EKS 6.11 releases

The following Amazon EMR 6.11.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.11.0-XXXX** release to view more details such as the related container image tag.

- [emr-6.11.0-latest](#)
- [emr-6.11.0-20230905](#)
- [emr-6.11.0-20230509](#)
- `emr-6.11.0-spark-rapids-latest`
- `emr-6.11.0-spark-rapids-20230509`
- `emr-6.11.0-java11-latest`
- `emr-6.11.0-java11-20230509`
- `notebook-spark/emr-6.11.0-latest`
- `notebook-spark/emr-6.11.0-20230509`
- `notebook-python/emr-6.11.0-latest`
- `notebook-python/emr-6.11.0-20230509`

Release notes

Release notes for Amazon EMR on EKS 6.11.0

- **Supported applications** - AWS SDK for Java 1.12.446, Apache Spark 3.3.2-amzn-0, Apache Hudi 0.13.0-amzn-0, Apache Iceberg 1.2.0-amzn-0, Delta 2.2.0, Apache Spark RAPIDS 23.02.0-amzn-0, Jupyter Enterprise Gateway 2.6.0
- **Supported components** - aws-sagemaker-spark-sdk, emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- **Supported configuration classifications**

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in the core-site.xml Hadoop file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the metrics.properties Spark file.
spark-defaults	Change values in the spark-defaults.conf Spark file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the hive-site.xml Spark file.
spark-log4j	Change values in the log4j.properties Spark file.

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway jupyter_enterprise_gateway_config.py file.

Classifications	Descriptions
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Notable features

The following features are included with the 6.11 release of Amazon EMR on EKS.

- **[Amazon EMR on EKS base image in Amazon ECR Public Gallery](#)** – If you use the [custom image](#) capability, our base image provides the essential jars, configuration, and libraries to interact with Amazon EMR on EKS. You can now find the base image in the [Amazon ECR Public Gallery](#).
- **[Spark container log rotation](#)** – Amazon EMR on EKS 6.11 supports Spark container log rotation. You can enable the capability with `containerLogRotationConfiguration` within the `MonitoringConfiguration` operation of the `StartJobRun` API. You can configure the `rotationSize` and `maxFilestoKeep` to specify the number and size of the log files that you want Amazon EMR on EKS to keep in the Spark driver and executor pods. For more information, see [Using Spark container log rotation](#).
- **[Volcano support in Spark operator and spark-submit](#)** – Amazon EMR on EKS 6.11 supports running Spark jobs with Volcano as Kubernetes custom scheduler in [Spark operator](#) and [spark-submit](#). You can use features like gang scheduling, queue management, preemption, and fair-share scheduling to achieve high scheduling throughput and optimized capacity. For more information, see [Using Volcano as a custom scheduler for Apache Spark on Amazon EMR on EKS](#).

emr-6.11.0-latest

Release notes: `emr-6.11.0-latest` currently points to `emr-20230905`.

Regions: `emr-6.11.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.11.0:latest`

emr-6.11.0-20230905

Release notes: 6.11.0-20230905 was released on September 29, 2023. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.11.0-20230509` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.11.0:20230509`

emr-6.11.0-20230509

Release notes: 6.11.0-20230509 was released on May 9, 2023. This is the initial release of Amazon EMR 6.11.0.

Regions: `emr-6.11.0-20230509` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.11.0:20230509`

Amazon EMR on EKS 6.10.0 releases

The following Amazon EMR 6.10.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.10.0-XXXX** release to view more details such as the related container image tag.

- [emr-6.10.0-latest](#)
- [emr-6.10.0-20230905](#)
- [emr-6.10.0-20230624](#)
- [emr-6.10.0-20230421](#)
- [emr-6.10.0-20230403](#)
- [emr-6.10.0-20230220](#)
- `emr-6.10.0-spark-rapids-latest`
- `emr-6.10.0-spark-rapids-20230624`
- `emr-6.10.0-spark-rapids-20230220`
- `emr-6.10.0-java11-latest`

- `emr-6.10.0-java11-20230624`
- `emr-6.10.0-java11-20230220`
- `notebook-spark/emr-6.10.0-latest`
- `notebook-spark/emr-6.10.0-20230624`
- `notebook-spark/emr-6.10.0-20230220`
- `notebook-python/emr-6.10.0-latest`
- `notebook-python/emr-6.10.0-20230624`
- `notebook-python/emr-6.10.0-20230220`

Release notes for Amazon EMR 6.10.0

- Supported applications - AWS SDK for Java 1.12.397, Spark 3.3.1-amzn-0, Hudi 0.12.2-amzn-0, Iceberg 1.1.0-amzn-0, Delta 2.2.0.
- Supported components - `aws-sagemaker-spark-sdk`, `emr-ddb`, `emr-goodies`, `emr-s3-select`, `emrfs`, `hadoop-client`, `hudi`, `hudi-spark`, `iceberg`, `spark-kubernetes`.
- Supported configuration classifications:

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
<code>core-site</code>	Change values in Hadoop's <code>core-site.xml</code> file.
<code>emrfs-site</code>	Change EMRFS settings.
<code>spark-metrics</code>	Change values in Spark's <code>metrics.properties</code> file.
<code>spark-defaults</code>	Change values in Spark's <code>spark-defaults.conf</code> file.
<code>spark-env</code>	Change values in the Spark environment.
<code>spark-hive-site</code>	Change values in Spark's <code>hive-site.xml</code> file.

Classifications	Descriptions
spark-log4j	Change values in Spark's log4j.properties file.

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway jupyter_enterprise_gateway_config.py file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as spark-hive-site.xml. For more information, see [Configure Applications](#).

Notable features

- **Spark operator** - With Amazon EMR on EKS 6.10.0 and higher, you can use the Kubernetes operator for Apache Spark, or *the Spark operator*, to deploy and manage Spark applications with the Amazon EMR release runtime on your own Amazon EKS clusters. For more information, see [Running Spark jobs with the Spark operator](#).
- **Java 11** - With Amazon EMR on EKS 6.10 and higher, you can launch Spark with Java 11 runtime. To do this, pass `emr-6.10.0-java11-latest` as a release label. We recommend that you validate and run performance tests before you move your production workloads from the Java 8 image to the Java 11 image.
- For the Amazon Redshift integration for Apache Spark, Amazon EMR on EKS 6.10.0 removes the dependency on `minimal-json.jar`, and automatically adds the required spark-redshift related jars to the executor class path for Spark: `spark-redshift.jar`, `spark-avro.jar`, and `RedshiftJDBC.jar`.

Changes

- EMRFS S3-optimized committer is now enabled by default for parquet, ORC, and text-based formats (including CSV and JSON).

emr-6.10.0-latest

Release notes: `emr-6.10.0-latest` currently points to `emr-6.10.0-20230905`.

Regions: `emr-6.10.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.10.0:latest`

emr-6.10.0-20230905

Release notes: `6.10.0-20230905` was released on September 29, 2023. Compared with the previous release, this version has been refreshed with recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.10.0-20230905` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.10.0:20230905`

emr-6.10.0-20230624

Release notes: `6.10.0-20230624` was released on July 7, 2023. Compared with the previous release, this version has been refreshed with recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.10.0-20230624` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.10.0:20230624`

emr-6.10.0-20230421

Release notes: `6.10.0-20230421` was released on April 28, 2023. Compared with the previous release, this version has been refreshed with recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.10.0-20230421` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.10.0:20230421`

emr-6.10.0-20230403

Release notes: `6.10.0-20230403` was released on April 12, 2023. Compared with the previous release, this version has been refreshed with recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.10.0-20230403` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.10.0:20230403`

emr-6.10.0-20230220

Release notes: `emr-6.10.0-20230220` was released on February 20, 2023. This is the initial release of Amazon EMR 6.10.0.

Regions: `emr-6.10.0-20230220` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.10.0:20230220`

Amazon EMR on EKS 6.9.0 releases

The following Amazon EMR 6.9.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.9.0-XXXX** release to view more details such as the related container image tag.

- [emr-6.9.0-latest](#)
- [emr-6.9.0-20230905](#)
- [emr-6.9.0-20230624](#)
- [emr-6.9.0-20221108](#)
- `emr-6.9.0-spark-rapids-latest`
- `emr-6.9.0-spark-rapids-20230624`
- `emr-6.9.0-spark-rapids-20221108`

- notebook-spark/emr-6.9.0-latest
- notebook-spark/emr-6.9.0-20230624
- notebook-spark/emr-6.9.0-20221108
- notebook-python/emr-6.9.0-latest
- notebook-python/emr-6.9.0-20230624
- notebook-python/emr-6.9.0-20221108

Release notes for Amazon EMR 6.9.0

- Supported applications - AWS SDK for Java 1.12.331, Spark 3.3.0-amzn-1, Hudi 0.12.1-amzn-0, Iceberg 0.14.1-amzn-0, Delta 2.1.0.
- Supported components - aws-sagemaker-spark-sdk, emr-ddb, emr-goodies, emr-s3-select, emrfs, hadoop-client, hudi, hudi-spark, iceberg, spark-kubernetes.
- Supported configuration classifications:

For use with [StartJobRun](#) and [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
core-site	Change values in Hadoop's core-site.xml file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in Spark's metrics.properties file.
spark-defaults	Change values in Spark's spark-defaults.conf file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in Spark's hive-site.xml file.
spark-log4j	Change values in Spark's log4j.properties file.

For use specifically with [CreateManagedEndpoint](#) APIs:

Classifications	Descriptions
jeg-config	Change values in Jupyter Enterprise Gateway <code>jupyter_enterprise_gateway_config.py</code> file.
jupyter-kernel-overrides	Change value for the Kernel Image in Jupyter Kernel Spec file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Notable features

- **Nvidia RAPIDS Accelerator for Apache Spark** - Amazon EMR on EKS to accelerate Spark using EC2 graphics processing unit (GPU) instance types. To use the Spark image with RAPIDS Accelerator, specify release label as `emr-6.9.0-spark-rapids-latest`. Visit the [documentation page](#) to learn more.
- **Spark-Redshift connector** - The Amazon Redshift integration for Apache Spark is included in Amazon EMR releases 6.9.0 and later. Previously an open-source tool, the native integration is a Spark connector that you can use to build Apache Spark applications that read from and write to data in Amazon Redshift and Amazon Redshift Serverless. For more information, see [Using Amazon Redshift integration for Apache Spark on Amazon EMR on EKS](#).
- **Delta Lake** - [Delta Lake](#) is an open-source storage format that enables building data lakes with transactional consistency, consistent definition of datasets, schema evolution changes, and data mutations support. Visit [Using Delta Lake](#) to learn more.
- **Modify PySpark parameters** - Interactive endpoints now support modifying Spark parameters associated with PySpark sessions in the EMR Studio Jupyter Notebook. Visit [Modifying PySpark session parameters](#) to learn more.

Resolved issues

- When you use the DynamoDB connector with Spark on Amazon EMR versions 6.6.0, 6.7.0, and 6.8.0, all reads from your table return an empty result, even though the input split references non-empty data. Amazon EMR release 6.9.0 fixes this issue.
- Amazon EMR on EKS 6.8.0 incorrectly populates the build hash in Parquet files metadata generated using [Apache Spark](#). This issue may cause tools that parse the metadata version string from Parquet files generated by Amazon EMR on EKS 6.8.0 to fail.

Known issue

- If you use the Amazon Redshift integration for Apache Spark and have a time, timetz, timestamp, or timestampz with microsecond precision in Parquet format, the connector rounds the time values to the nearest millisecond value. As a workaround, use the text unload format `unload_s3_format` parameter.

emr-6.9.0-latest

Release notes: `emr-6.9.0-latest` currently points to `emr-6.9.0-20230905`.

Regions: `emr-6.9.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.9.0:latest`

emr-6.9.0-20230905

Release notes: `emr-6.9.0-20230905`. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.9.0-20230905` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.9.0:20230905`

emr-6.9.0-20230624

Release notes: `emr-6.9.0-20230624` was released on July 7, 2023.

Regions: `emr-6.9.0-20230624` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.9.0:20230624`

emr-6.9.0-20221108

Release notes: `emr-6.9.0-20221108` was released on December 08, 2022. This is the initial release of Amazon EMR 6.9.0.

Regions: `emr-6.9.0-20221108` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.9.0:20221108`

Amazon EMR on EKS 6.8.0 releases

The following Amazon EMR 6.8.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.8.0-XXXX** release to view more details such as the related container image tag.

- [emr-6.8.0-latest](#)
- [emr-6.8.0-20230905](#)
- [emr-6.8.0-20230624](#)
- [emr-6.8.0-20221219](#)
- [emr-6.8.0-20220802](#)

Release notes for Amazon EMR 6.8.0

- Supported applications - AWS SDK for Java 1.12.170, Spark 3.3.0-amzn-0, Hudi 0.11.1-amzn-0, Iceberg 0.14.0-amzn-0.
- Supported components - `aws-sagemaker-spark-sdk`, `emr-ddb`, `emr-goodies`, `emr-s3-select`, `emrfs`, `hadoop-client`, `hudi`, `hudi-spark`, `iceberg`, `spark-kubernetes`.
- Supported configuration classifications:

Classifications	Descriptions
<code>core-site</code>	Change values in Hadoop's <code>core-site.xml</code> file.
<code>emrfs-site</code>	Change EMRFS settings.

Classifications	Descriptions
spark-metrics	Change values in Spark's metrics.properties file.
spark-defaults	Change values in Spark's spark-defaults.conf file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in Spark's hive-site.xml file.
spark-log4j	Change values in Spark's log4j.properties file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configure Applications](#).

Notable features

- **Spark3.3.0** - Amazon EMR on EKS 6.8 includes Spark 3.3.0, which supports using separate node selector labels for Spark driver executor pods. These new labels enable you to define the node types for the driver and executor pods separately in the StartJobRun API, without using pod templates.
 - Driver node selector property: `spark.kubernetes.driver.node.selector.[labelKey]`
 - Executor node selector property: `spark.kubernetes.executor.node.selector.[labelKey]`
- **Enhanced job failure message** - This release introduces the configuration `spark.stage.extraDetailsOnFetchFailures.enabled` and `spark.stage.extraDetailsOnFetchFailures.maxFailuresToInclude` to track task failures due to user code. These details will be used to enhance the failure message displayed in the driver log when a stage is aborted due to shuffle fetch failure.

Property name	Default value	Meaning	Since version
<code>spark.stage.extraD</code>	false	If set to true, this property is used	emr-6.8

Property name	Default value	Meaning	Since version
<code>etailsOnF etchFailu res.enabled</code>		to enhance the job failure message displayed in the driver log when a stage is aborted due to Shuffle Fetch Failures. By default the last 5 task failures caused by user code is tracked, and the failure error message is appended in the Driver Logs. To increase the number of task failures with user exceptions to track, see the config <code>spark.stage.extraDetailsOnFetchFailures.maxFailuresToInclude</code>.	

Property name	Default value	Meaning	Since version
<code>spark.stage.extraDetailsOnFetchFailures.maxFailuresToInclude</code>	5	Number of task failures to track per stage and attempt. This property is used to enhance the job failure message with user exceptions displayed in the driver log when a stage is aborted due to Shuffle Fetch Failures. This property works only if Config <code>spark.stage.extraDetailsOnFetchFailures.enabled</code> is set to true.	emr-6.8

For more information see the [Apache Spark configuration documentation](#).

Known issue

- Amazon EMR on EKS 6.8.0 incorrectly populates the build hash in Parquet files metadata generated using [Apache Spark](#). This issue may cause tools that parse the metadata version string from Parquet files generated by Amazon EMR on EKS 6.8.0 to fail. Customers who parse the version string from Parquet metadata and depend on build hash should switch to a different Amazon EMR version and rewrite the file.

Resolved issue

- Interrupt Kernel capability for pySpark kernels** - In progress interactive workloads that are triggered by executing cells in a notebook can be stopped by using the Interrupt Kernel capability. A fix has been introduced so that this functionality works for pySpark kernels. This is

also available in open source at [Changes for handling interrupts for PySpark Kubernetes Kernel #1115](#).

emr-6.8.0-latest

Release notes: `emr-6.8.0-latest` currently points to `emr-6.8.0-20230624`.

Regions: `emr-6.8.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.8.0:latest`

emr-6.8.0-20230905

Release notes: `emr-6.8.0-20230905` was released on September 29, 2023. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.8.0-20230905` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.8.0:20230905`

emr-6.8.0-20230624

Release notes: `emr-6.8.0-20230624` was released on July 7, 2023. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.8.0-20230624` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.8.0:20230624`

emr-6.8.0-20221219

Release notes: `emr-6.8.0-20221219` was released on Jan 19, 2023. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.8.0-20221219` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.8.0:20221219`

emr-6.8.0-20220802

Release notes: `emr-6.8.0-20220802` was released on Sep 27, 2022. This is the initial release of Amazon EMR 6.8.0.

Regions: `emr-6.8.0-20220802` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.8.0:20220802`

Amazon EMR on EKS 6.7.0 releases

The following Amazon EMR 6.7.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.7.0-XXXX** release to view more details such as the related container image tag.

- [emr-6.7.0-latest](#)
- [emr-6.7.0-20240321](#)
- [emr-6.7.0-20230624](#)
- [emr-6.7.0-20221219](#)
- [emr-6.7.0-20220630](#)

Release notes for Amazon EMR 6.7.0

- Supported applications - Spark 3.2.1-amzn-0, Jupyter Enterprise Gateway 2.6, Hudi 0.11-amzn-0, Iceberg 0.13.1.
- Supported components - `aws-hm-client` (Glue connector), `aws-sagemaker-spark-sdk`, `emr-s3-select`, `emrfs`, `emr-ddb`, `hudi-spark`.
- With the upgrade to JEG 2.6, kernel management is now asynchronous, which means that JEG does not block transactions when a kernel launch is in progress. This greatly improves the user experience by providing the following:
 - capability to execute commands in currently running notebooks when other kernel launches are in progress

- capability to launch multiple kernels simultaneously without impacting already running kernels
- Supported configuration classifications:

Classifications	Descriptions
core-site	Change values in the Hadoop core-site.xml file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in the Spark metrics.properties file.
spark-defaults	Change values in the Spark spark-defaults.conf file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in the Spark hive-site.xml file.
spark-log4j	Change values in the Spark log4j.properties file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configuring Applications](#).

Resolved issues

- Amazon EMR on EKS 6.7 fixes an issue in 6.6 when using Apache Spark's pod templates functionality with interactive endpoints. The issue was present in Amazon EMR on EKS releases 6.4, 6.5 and 6.6. You can now use pod templates to define how your Spark driver and executor pods start when using interactive endpoints to run interactive analytics.
- In previous Amazon EMR on EKS releases, Jupyter Enterprise Gateway would block transactions when kernel launch was in progress, and this impeded the execution of currently running

notebook sessions. You can now execute commands in currently running notebooks when other kernel launches are in progress. You can also launch multiple kernels simultaneously without the risk of losing connectivity to kernels that are already running.

emr-6.7.0-latest

Release notes: `emr-6.7.0-latest` currently points to `emr-6.7.0-20240321`.

Regions: `emr-6.7.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.7.0:latest`

emr-6.7.0-20240321

Release notes: `emr-6.7.0-20240321` was released on March 11, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.7.0-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.7.0:20240321`

emr-6.7.0-20230624

Release notes: `emr-6.7.0-20230624` was released on July 7, 2023. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.7.0-20230624` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.7.0:20230624`

emr-6.7.0-20221219

Release notes: `emr-6.7.0-20221219` was released on Jan. 19, 2023. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.7.0-20221219` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.7.0:20221219`

emr-6.7.0-20220630

Release notes: `emr-6.7.0-20220630` was released on July 12, 2022. This is the initial release of Amazon EMR 6.7.0.

Regions: `emr-6.7.0-20220630` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.7.0:20220630`

Amazon EMR on EKS 6.6.0 releases

The following Amazon EMR 6.6.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.6.0-XXXX** release to view more details such as the related container image tag.

- [emr-6.6.0-latest](#)
- [emr-6.6.0-20240321](#)
- [emr-6.6.0-20230624](#)
- [emr-6.6.0-20221219](#)
- [emr-6.6.0-20220411](#)

Release notes for Amazon EMR 6.6.0

- Supported applications - Spark 3.2.0-amzn-0, Jupyter Enterprise Gateway (endpoints, public preview), Hudi 0.10.1-amzn-0, Iceberg 0.13.1.
- Supported components - `aws-hm-client` (Glue connector), `aws-sagemaker-spark-sdk`, `emr-s3-select`, `emrfs`, `emr-ddb`, `hudi-spark`.
- Supported configuration classifications:

Classifications	Descriptions
<code>core-site</code>	Change values in Hadoop's <code>core-site.xml</code> file.

Classifications	Descriptions
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in Spark's metrics.properties file.
spark-defaults	Change values in Spark's spark-defaults.conf file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in Spark's hive-site.xml file.
spark-log4j	Change values in Spark's log4j.properties file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as spark-hive-site.xml. For more information, see [Configuring Applications](#).

Known issue

- Spark pod template functionality with interactive endpoints is not working in Amazon EMR on EKS release 6.4, 6.5, and 6.6.

Resolved issue

- Interactive endpoint logs are uploaded to Cloudwatch and S3.

emr-6.6.0-latest

Release notes: `emr-6.6.0-latest` currently points to `emr-6.6.0-20240321`.

Regions: `emr-6.6.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.6.0:latest`

emr-6.6.0-20240321

Release notes: `emr-6.6.0-20240321` was released on March 11, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.6.0-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.6.0:20240321`

emr-6.6.0-20230624

Release notes: `emr-6.6.0-20230624` was released on Jan 27, 2023. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.6.0-20230624` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.6.0:20230624`

emr-6.6.0-20221219

Release notes: `emr-6.6.0-20221219` was released on Jan 27, 2023. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.6.0-20221219` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.6.0:20221219`

emr-6.6.0-20220411

Release notes: `emr-6.6.0-20220411` was released on May 20, 2022. This is the initial release of Amazon EMR 6.6.0.

Regions: `emr-6.6.0-20220411` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.6.0:20220411`

Amazon EMR on EKS 6.5.0 releases

The following Amazon EMR 6.5.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.5.0-XXXX** release to view more details such as the related container image tag.

- [emr-6.5.0-latest](#)
- [emr-6.5.0-20240321](#)
- [emr-6.5.0-20221219](#)
- [emr-6.5.0-20220802](#)
- [emr-6.5.0-20211119](#)

Release notes for Amazon EMR 6.5.0

- Supported applications - Spark 3.1.2-amzn-1, Jupyter Enterprise Gateway (endpoints, public preview).
- Supported components - `aws-hm-client` (Glue connector), `aws-sagemaker-spark-sdk`, `emr-s3-select`, `emrfs`, `emr-ddb`, `hudi-spark`.
- Supported configuration classifications:

Classifications	Descriptions
<code>core-site</code>	Change values in Hadoop's <code>core-site.xml</code> file.
<code>emrfs-site</code>	Change EMRFS settings.
<code>spark-metrics</code>	Change values in Spark's <code>metrics.properties</code> file.
<code>spark-defaults</code>	Change values in Spark's <code>spark-defaults.conf</code> file.
<code>spark-env</code>	Change values in the Spark environment.
<code>spark-hive-site</code>	Change values in Spark's <code>hive-site.xml</code> file.

Classifications	Descriptions
spark-log4j	Change values in Spark's log4j.properties file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as spark-hive-site.xml. For more information, see [Configuring Applications](#).

Known Issue

- Spark pod template functionality with interactive endpoints is not working in Amazon EMR on EKS releases 6.4 and 6.5.

emr-6.5.0-latest

Release notes: `emr-6.5.0-latest` currently points to `emr-6.5.0-20240321`.

Regions: `emr-6.5.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.5.0:latest`

emr-6.5.0-20240321

Release notes: `emr-6.5.0-20240321` was released on March 11, 2024. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.5.0-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.5.0:20240321`

emr-6.5.0-20221219

Release notes: `emr-6.5.0-20221219` was released on Jan 19, 2023. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.5.0-20221219` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.5.0:20221219`

emr-6.5.0-20220802

Release notes: `emr-6.5.0-20220802` was released on Aug 24, 2022. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages.

Regions: `emr-6.5.0-20220802` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.5.0:20220802`

emr-6.5.0-20211119

Release notes: `emr-6.5.0-20211119` was released on Jan 20, 2022. This is the initial release of Amazon EMR 6.5.0.

Regions: `emr-6.5.0-20211119` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.5.0:20211119`

Amazon EMR on EKS 6.4.0 releases

The following Amazon EMR 6.4.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.4.0-XXXX** release to view more details such as the related container image tag.

- [emr-6.4.0-latest](#)
- [emr-6.4.0-20240321](#)
- [emr-6.4.0-20221219](#)
- [emr-6.4.0-20210830](#)

Release notes for Amazon EMR 6.4.0

- Supported applications - Spark 3.1.2-amzn-0, Jupyter Enterprise Gateway (endpoints, public preview).

- Supported components - `aws-hm-client` (Glue connector), `aws-sagemaker-spark-sdk`, `emr-s3-select`, `emrfs`, `emr-ddb`, `hudi-spark`.
- Supported configuration classifications:

Classifications	Descriptions
<code>core-site</code>	Change values in Hadoop's <code>core-site.xml</code> file.
<code>emrfs-site</code>	Change EMRFS settings.
<code>spark-metrics</code>	Change values in Spark's <code>metrics.properties</code> file.
<code>spark-defaults</code>	Change values in Spark's <code>spark-defaults.conf</code> file.
<code>spark-env</code>	Change values in the Spark environment.
<code>spark-hive-site</code>	Change values in Spark's <code>hive-site.xml</code> file.
<code>spark-log4j</code>	Change values in Spark's <code>log4j.properties</code> file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as `spark-hive-site.xml`. For more information, see [Configuring Applications](#).

Known issue

- Spark pod template functionality with interactive endpoints is not working in Amazon EMR on EKS release 6.4.

emr-6.4.0-latest

Release notes: `emr-6.4.0-latest` currently points to `emr-6.4.0-20240321`.

Regions: `emr-6.4.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.4.0:latest`

emr-6.4.0-20240321

Release notes: `emr-6.4.0-20240321` was released on March 11, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.4.0-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.4.0:20240321`

emr-6.4.0-20221219

Release notes: `emr-6.4.0-20221219` was released on Jan 27, 2023. Compared to the previous version, this version has been refreshed with the recently added Amazon Linux packages.

Regions: `emr-6.4.0-20221219` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.4.0:20221219`

emr-6.4.0-20210830

Release notes: `emr-6.4.0-20210830` was released on Dec 9, 2021. This is the initial release of Amazon EMR 6.4.0.

Regions: `emr-6.4.0-20210830` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.4.0:20210830`

Amazon EMR on EKS 6.3.0 releases

The following Amazon EMR 6.3.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.3.0-XXXX** release to view more details such as the related container image tag.

- [emr-6.3.0-latest](#)

- [emr-6.3.0-20240321](#)
- [emr-6.3.0-20220802](#)
- [emr-6.3.0-20211008](#)
- [emr-6.3.0-20210802](#)
- [emr-6.3.0-20210429](#)

Release notes for Amazon EMR 6.3.0

- New features - Beginning with Amazon EMR 6.3.0 in the 6.x release series, Amazon EMR on EKS supports Spark's pod template feature. You can also turn on the Spark event log rotation feature for Amazon EMR on EKS. For more information, see [Using pod templates](#) and [Using Spark event log rotation](#).
- Supported applications - Spark 3.1.1-amzn-0, Jupyter Enterprise Gateway (endpoints, public preview).
- Supported components - aws-hm-client (Glue connector), aws-sagemaker-spark-sdk, emr-s3-select, emrfs, emr-ddb, hudi-spark.
- Supported configuration classifications:

Classifications	Descriptions
core-site	Change values in Hadoop's core-site.xml file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in Spark's metrics.properties file.
spark-defaults	Change values in Spark's spark-defaults.conf file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in Spark's hive-site.xml file.
spark-log4j	Change values in Spark's log4j.properties file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as spark-hive-site.xml. For more information, see [Configuring Applications](#).

emr-6.3.0-latest

Release notes: `emr-6.3.0-latest` currently points to `emr-6.3.0-20240321`.

Regions: `emr-6.3.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.3.0:latest`

emr-6.3.0-20240321

Release notes: `emr-6.3.0-20240321` was released on March 11, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.3.0-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.3.0:20240321`

emr-6.3.0-20220802

Release notes: `emr-6.3.0-20220802` was released on Sep 27, 2022. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages.

Regions: `emr-6.3.0-20220802` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.3.0:20220802`

emr-6.3.0-20211008

Release notes: `emr-6.3.0-20211008` was released on Dec 9, 2021. Compared to the previous version, this version contains issue fixes and security updates.

Regions: `emr-6.3.0-20211008` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.3.0:20211008`

emr-6.3.0-20210802

Release notes: `emr-6.3.0-20210802` was released on Aug 2, 2021. Compared to the previous version, this version contains issue fixes and security updates.

Regions: `emr-6.3.0-20210802` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.3.0:20210802`

emr-6.3.0-20210429

Release notes: `emr-6.3.0-20210429` was released on April 29, 2021. This is the initial release of Amazon EMR 6.3.0.

Regions: `emr-6.3.0-20210429` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.3.0:20210429`

Amazon EMR on EKS 6.2.0 releases

The following Amazon EMR 6.2.0 releases are available for Amazon EMR on EKS. Select a specific **emr-6.2.0-XXXX** release to view more details such as the related container image tag.

- [emr-6.2.0-latest](#)
- [emr-6.2.0-20240321](#)
- [emr-6.2.0-20220802](#)
- [emr-6.2.0-20211008](#)
- [emr-6.2.0-20210802](#)
- [emr-6.2.0-20210615](#)
- [emr-6.2.0-20210129](#)
- [emr-6.2.0-20201218](#)
- [emr-6.2.0-20201201](#)

Release notes for Amazon EMR 6.2.0

- Supported applications - Spark 3.0.1-amzn-0, Jupyter Enterprise Gateway (endpoints, public preview).
- Supported components - aws-hm-client (Glue connector), aws-sagemaker-spark-sdk, emr-s3-select, emrfs, emr-ddb, hudi-spark.
- Supported configuration classifications:

Classifications	Descriptions
core-site	Change values in Hadoop's core-site.xml file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in Spark's metrics.properties file.
spark-defaults	Change values in Spark's spark-defaults.conf file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in Spark's hive-site.xml file.
spark-log4j	Change values in Spark's log4j.properties file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as spark-hive-site.xml. For more information, see [Configuring Applications](#).

emr-6.2.0-latest

Release notes: emr-6.2.0-latest currently points to emr-6.2.0-20240321.

Regions: emr-6.2.0-latest is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: emr-6.2.0:20240321

emr-6.2.0-20240321

Release notes: `emr-6.2.0-20240321` was released on March 11, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-6.2.0-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.2.0:20240321`

emr-6.2.0-20220802

Release notes: `emr-6.2.0-20220802` was released on Sep 27, 2022. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages.

Regions: `emr-6.2.0-20220802` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-6.2.0:20220802`

emr-6.2.0-20211008

Release notes: `emr-6.2.0-20211008` was released on Dec 9, 2021. Compared to the previous version, this version contains issue fixes and security updates.

Regions: `emr-6.2.0-20211008` is available in the following Regions: US East (N. Virginia), US West (Oregon), Asia Pacific (Tokyo), Europe (Ireland), South America (Sao Paulo).

Container image tag: `emr-6.2.0:20211008`

emr-6.2.0-20210802

Release notes: `emr-6.2.0-20210802` was released on Aug 2, 2021. Compared to the previous version, this version contains issue fixes and security updates.

Regions: `emr-6.2.0-20210802` is available in the following Regions: US East (N. Virginia), US West (Oregon), Asia Pacific (Tokyo), Europe (Ireland), South America (Sao Paulo).

Container image tag: `emr-6.2.0:20210802`

emr-6.2.0-20210615

Release notes: `emr-6.2.0-20210615` was released on June 15, 2021. Compared to the previous version, this version contains issue fixes and security updates.

Regions: `emr-6.2.0-20210615` is available in the following Regions: US East (N. Virginia), US West (Oregon), Asia Pacific (Tokyo), Europe (Ireland), South America (Sao Paulo).

Container image tag: `emr-6.2.0:20210615`

emr-6.2.0-20210129

Release notes: `emr-6.2.0-20210129` was released on January 29, 2021. Compared to `emr-6.2.0-20201218`, this version contains issue fixes and security updates.

Regions: `emr-6.2.0-20210129` is available in the following Regions: US East (N. Virginia), US West (Oregon), Asia Pacific (Tokyo), Europe (Ireland), South America (Sao Paulo).

Container image tag: `emr-6.2.0-20210129`

emr-6.2.0-20201218

Release notes: `emr-6.2.0-20201218` was released on December 18, 2020. Compared to `emr-6.2.0-20201201`, this version contains issue fixes and security updates.

Regions: `emr-6.2.0-20201218` is available in the following Regions: US East (N. Virginia), US West (Oregon), Asia Pacific (Tokyo), Europe (Ireland), South America (Sao Paulo).

Container image tag: `emr-6.2.0-20201218`

emr-6.2.0-20201201

Release notes: `emr-6.2.0-20201201` was released on December 1, 2020. This is the initial release of Amazon EMR 6.2.0.

Regions: `emr-6.2.0-20201201` is available in the following Regions: US East (N. Virginia), US West (Oregon), Asia Pacific (Tokyo), Europe (Ireland), South America (Sao Paulo).

Container image tag: `emr-6.2.0-20201201`

Amazon EMR on EKS 5.36.0 releases

The following Amazon EMR 5.36.0 releases are available for Amazon EMR on EKS. Select a specific **emr-5.36.0-XXXX** release to view more details such as the related container image tag.

- [emr-5.36.0-latest](#)
- [emr-5.36.0-20240321](#)
- [emr-5.36.0-20221219](#)
- [emr-5.36.0-20220620](#)
- [emr-5.36.0-20220525](#)

Release notes for Amazon EMR 5.36.0

- Fixed log4j2 security issues.
- Supported applications - Spark 2.4.8-amzn-2, Jupyter Enterprise Gateway (endpoints, public preview; Scala kernel is not supported), livy-0.7.1, fluentd-4.0.0.
- Supported components - aws-hm-client, aws-sagemaker-spark-sdk, emr-ddb, emr-goodies, emr-kinesis, kerberos-server.
- Supported configuration classifications:

Classifications	Descriptions
core-site	Change values in Hadoop's core-site.xml file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in Spark's metrics.properties file.
spark-defaults	Change values in Spark's spark-defaults.conf file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in Spark's hive-site.xml file.
spark-log4j	Change values in Spark's log4j.properties file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as spark-hive-site.xml. For more information, see [Configuring Applications](#).

emr-5.36.0-latest

Release notes: `emr-5.36.0-latest` currently points to `emr-5.36.0-20240321`.

Regions: `emr-5.36.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.36.0:latest`

emr-5.36.0-20240321

Release notes: `emr-5.36.0-20240321` was released on March 11, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-5.36.0-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.36.0:20240321`

emr-5.36.0-20221219

Release notes: `emr-5.36.0-20221219` was released on Jan 27, 2023. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages.

Regions: `emr-5.36.0-20221219` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.36.0:20221219`

emr-5.36.0-20220620

Release notes: `emr-5.36.0-20220620` was released on July 27, 2022. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages.

Regions: `emr-5.36.0-20220620` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.36.0:20220620`

emr-5.36.0-20220525

Release notes: `emr-5.36.0-20220525` was released on June 16, 2022. This is the initial release of Amazon EMR 5.36.0.

Regions: `emr-5.36.0-20220525` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.36.0:20220525`

Amazon EMR on EKS 5.35.0 releases

The following Amazon EMR 5.35.0 releases are available for Amazon EMR on EKS. Select a specific **emr-5.35.0-XXXX** release to view more details such as the related container image tag.

- [emr-5.35.0-latest](#)
- [emr-5.35.0-20240321](#)
- [emr-5.35.0-20221219](#)
- [emr-5.35.0-20220802](#)
- [emr-5.35.0-20220307](#)

Release notes for Amazon EMR 5.35.0

- Fixed log4j2 security issues.
- Supported applications - Spark 2.4.8-amzn-1, Hudi 0.9.0-amzn-2, Jupyter Enterprise Gateway (endpoints, public preview; Scala kernel is not supported).
- Supported components - aws-hm-client (Glue connector), aws-sagemaker-spark-sdk, emr-s3-select, emrfs, emr-ddb, hudi-spark.
- Supported configuration classifications:

Classifications	Descriptions
core-site	Change values in Hadoop's core-site.xml file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in Spark's metrics.properties file.
spark-defaults	Change values in Spark's spark-defaults.conf file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in Spark's hive-site.xml file.
spark-log4j	Change values in Spark's log4j.properties file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as spark-hive-site.xml. For more information, see [Configuring Applications](#).

emr-5.35.0-latest

Release notes: `emr-5.35.0-latest` currently points to `emr-5.35.0-20240321`.

Regions: `emr-5.35.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.35.0:latest`

emr-5.35.0-20240321

Release notes: `emr-5.35.0-20240321` was released on March 11, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-5.35.0-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.35.0:20240321`

emr-5.35.0-20221219

Release notes: `emr-5.35.0-20221219` was released on Jan 27, 2023. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages.

Regions: `emr-5.35.0-20221219` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.35.0:20221219`

emr-5.35.0-20220802

Release notes: `emr-5.35.0-20220802` was released on Sep 27, 2022. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages.

Regions: `emr-5.35.0-20220802` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.35.0:20220802`

emr-5.35.0-20220307

Release notes: `emr-5.35.0-20220307` was released on Mar 30, 2022. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages.

Regions: `emr-5.35.0-20220307` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.35.0:20220307`

Amazon EMR on EKS 5.34.0 releases

The following Amazon EMR 5.34.0 releases are available for Amazon EMR on EKS. Select a specific **emr-5.34.0-XXXX** release to view more details such as the related container image tag.

- [emr-5.34.0-latest](#)

- [emr-5.34.0-20240321](#)
- [emr-5.34.0-20220802](#)

Release notes for Amazon EMR 5.34.0

- Supported applications - Spark 2.4.8-amzn-0, Jupyter Enterprise Gateway (endpoints, public preview; Scala kernel is not supported).
- Supported components - aws-hm-client (Glue connector), aws-sagemaker-spark-sdk, emr-s3-select, emrfs, emr-ddb, hudi-spark.
- Supported configuration classifications:

Classifications	Descriptions
core-site	Change values in Hadoop's core-site.xml file.
emrfs-site	Change EMRFS settings.
spark-metrics	Change values in Spark's metrics.properties file.
spark-defaults	Change values in Spark's spark-defaults.conf file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in Spark's hive-site.xml file.
spark-log4j	Change values in Spark's log4j.properties file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as spark-hive-site.xml. For more information, see [Configuring Applications](#).

emr-5.34.0-latest

Release notes: `emr-5.34.0-latest` currently points to `emr-5.34.0-20220802`.

Regions: `emr-5.34.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.34.0:latest`

emr-5.34.0-20240321

Release notes: `emr-5.34.0-20240321` was released on March 11, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-5.34.0-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.34.0:20240321`

emr-5.34.0-20220802

Release notes: `emr-5.34.0-20220802` was released on Aug 24, 2022. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages.

Regions: `emr-5.34.0-20220802` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.34.0:20220802`

emr-5.34.0-20211208

Release notes: `emr-5.34.0-20211208` was released on Jan 20, 2022. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages.

Regions: `emr-5.34.0-20211208` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.34.0:20211208`

Amazon EMR on EKS 5.33.0 releases

The following Amazon EMR 5.33.0 releases are available for Amazon EMR on EKS. Select a specific **emr-5.33.0-XXXX** release to view more details such as the related container image tag.

- [emr-5.33.0-latest](#)
- [emr-5.33.0-20240321](#)
- [emr-5.33.0-20221219](#)
- [emr-5.33.0-20220802](#)
- [emr-5.33.0-20211008](#)
- [emr-5.33.0-20210802](#)
- [emr-5.33.0-20210615](#)
- [emr-5.33.0-20210323](#)

Release notes for Amazon EMR 5.33.0

- New feature - Beginning with Amazon EMR 5.33.0 in the 5.x release series, Amazon EMR on EKS supports Spark's pod template feature. For more information, see [Using pod templates](#).
- Supported applications - Spark 2.4.7-amzn-1, Jupyter Enterprise Gateway (endpoints, public preview; Scala kernel is not supported).
- Supported components - `aws-hm-client` (Glue connector), `aws-sagemaker-spark-sdk`, `emr-s3-select`, `emrfs`, `emr-ddb`, `hudi-spark`.
- Supported configuration classifications:

Classifications	Descriptions
<code>core-site</code>	Change values in Hadoop's <code>core-site.xml</code> file.
<code>emrfs-site</code>	Change EMRFS settings.
<code>spark-metrics</code>	Change values in Spark's <code>metrics.properties</code> file.
<code>spark-defaults</code>	Change values in Spark's <code>spark-defaults.conf</code> file.
<code>spark-env</code>	Change values in the Spark environment.
<code>spark-hive-site</code>	Change values in Spark's <code>hive-site.xml</code> file.
<code>spark-log4j</code>	Change values in Spark's <code>log4j.properties</code> file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as spark-hive-site.xml. For more information, see [Configuring Applications](#).

emr-5.33.0-latest

Release notes: `emr-5.33.0-latest` currently points to `emr-5.33.0-20240321`.

Regions: `emr-5.33.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.33.0:latest`

emr-5.33.0-20240321

Release notes: `emr-5.33.0-20240321` was released on March 11, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-5.33.0-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.33.0:20240321`

emr-5.33.0-20221219

Release notes: `emr-5.33.0-20221219` was released on Jan 19, 2023. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-5.33.0-20221219` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.33.0:20221219`

emr-5.33.0-20220802

Release notes: `emr-5.33.0-20220802` was released on Aug 24, 2022. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages.

Regions: `emr-5.33.0-20220802` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.33.0:20220802`

emr-5.33.0-20211008

Release notes: `emr-5.33.0-20211008` was released on Dec 9, 2021. Compared to the previous version, this version contains issue fixes and security updates.

Regions: `emr-5.33.0-20211008` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.33.0:20211008`

emr-5.33.0-20210802

Release notes: `emr-5.33.0-20210802` was released on Aug 2, 2021. Compared to the previous version, this version contains issue fixes and security updates.

Regions: `emr-5.33.0-20210802` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.33.0:20210802`

emr-5.33.0-20210615

Release notes: `emr-5.33.0-20210615` was released on June 15, 2021. Compared to the previous version, this version contains issue fixes and security updates.

Regions: `emr-5.33.0-20210615` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.33.0:20210615`

emr-5.33.0-20210323

Release notes: `emr-5.33.0-20210323` was released on March 23, 2021. This is the initial release of Amazon EMR 5.33.0.

Regions: `emr-5.33.0-20210323` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.33.0-20210323`

Amazon EMR on EKS 5.32.0 releases

The following Amazon EMR 5.32.0 releases are available for Amazon EMR on EKS. Select a specific **emr-5.32.0-XXXX** release to view more details such as the related container image tag.

- [emr-5.32.0-latest](#)
- [emr-5.32.0-20240321](#)
- [emr-5.32.0-20220802](#)
- [emr-5.32.0-20211008](#)
- [emr-5.32.0-20210802](#)
- [emr-5.32.0-20210615](#)
- [emr-5.32.0-20210129](#)
- [emr-5.32.0-20201218](#)
- [emr-5.32.0-20201201](#)

Release notes for Amazon EMR 5.32.0

- Supported applications - Spark 2.4.7-amzn-0, Jupyter Enterprise Gateway (endpoints, public preview; Scala kernel is not supported).
- Supported components - `aws-hm-client` (Glue connector), `aws-sagemaker-spark-sdk`, `emr-s3-select`, `emrfs`, `emr-ddb`, `hudi-spark`.
- Supported configuration classifications:

Classifications	Descriptions
<code>core-site</code>	Change values in Hadoop's <code>core-site.xml</code> file.
<code>emrfs-site</code>	Change EMRFS settings.
<code>spark-metrics</code>	Change values in Spark's <code>metrics.properties</code> file.

Classifications	Descriptions
spark-defaults	Change values in Spark's spark-defaults.conf file.
spark-env	Change values in the Spark environment.
spark-hive-site	Change values in Spark's hive-site.xml file.
spark-log4j	Change values in Spark's log4j.properties file.

Configuration classifications allow you to customize applications. These often correspond to a configuration XML file for the application, such as spark-hive-site.xml. For more information, see [Configuring Applications](#).

emr-5.32.0-latest

Release notes: `emr-5.32.0-latest` currently points to `emr-5.32.0-20240321`.

Regions: `emr-5.32.0-latest` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.32.0:latest`

emr-5.32.0-20240321

Release notes: `emr-5.32.0-20240321` was released on March 11, 2024. Compared to the previous release, this release has been refreshed with the recently updated Amazon Linux packages and critical fixes.

Regions: `emr-5.32.0-20240321` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.32.0:20240321`

emr-5.32.0-20220802

Release notes: `emr-5.32.0-20220802` was released on Aug 24, 2022. Compared to the previous version, this version has been refreshed with the recently updated Amazon Linux packages.

Regions: `emr-5.32.0-20220802` is available in all Regions supported by Amazon EMR on EKS. For more information, see [Amazon EMR on EKS service endpoints](#).

Container image tag: `emr-5.32.0:20220802`

emr-5.32.0-20211008

Release notes: `emr-5.32.0-20211008` was released on Dec 9, 2021. Compared to the previous version, this version contains issue fixes and security updates.

Regions: `emr-5.32.0-20211008` is available in the following Regions: US East (N. Virginia), US West (Oregon), Asia Pacific (Tokyo), Europe (Ireland), South America (Sao Paulo).

Container image tag: `emr-5.32.0:20211008`

emr-5.32.0-20210802

Release notes: `emr-5.32.0-20210802` was released on Aug 2, 2021. Compared to the previous version, this version contains issue fixes and security updates.

Regions: `emr-5.32.0-20210802` is available in the following Regions: US East (N. Virginia), US West (Oregon), Asia Pacific (Tokyo), Europe (Ireland), South America (Sao Paulo).

Container image tag: `emr-5.32.0:20210802`

emr-5.32.0-20210615

Release notes: `emr-5.32.0-20210615` was released on June 15, 2021. Compared to the previous version, this version contains issue fixes and security updates.

Regions: `emr-5.32.0-20210615` is available in the following Regions: US East (N. Virginia), US West (Oregon), Asia Pacific (Tokyo), Europe (Ireland), South America (Sao Paulo).

Container image tag: `emr-5.32.0:20210615`

emr-5.32.0-20210129

Release notes: `emr-5.32.0-20210129` was released on January 29, 2021. Compared to `emr-5.32.0-20201218`, this version contains issue fixes and security updates.

Regions: `emr-5.32.0-20210129` is available in the following Regions: US East (N. Virginia), US West (Oregon), Asia Pacific (Tokyo), Europe (Ireland), South America (Sao Paulo).

Container image tag: `emr-5.32.0-20210129`

emr-5.32.0-20201218

Release notes: 5.32.0-20201218 was released on December 18, 2020. Compared to 5.32.0-20201201, this version contains issue fixes and security updates.

Regions: `emr-5.32.0-20201218` is available in the following Regions: US East (N. Virginia), US West (Oregon), Asia Pacific (Tokyo), Europe (Ireland), South America (Sao Paulo).

Container image tag: `emr-5.32.0-20201218`

emr-5.32.0-20201201

Release notes: 5.32.0-20201201 was released on December 1, 2020. This is the initial release of Amazon EMR 5.32.0.

Regions: 5.32.0-20201201d is available in the following Regions: US East (N. Virginia), US West (Oregon), Asia Pacific (Tokyo), Europe (Ireland), South America (Sao Paulo).

Container image tag: `emr-5.32.0-20201201`

Document history

The following table describes the important changes to the documentation since the last release of Amazon EMR on EKS. For more information about updates to this documentation, you can subscribe to an RSS feed.

Change	Description	Date
Added emr-spark-8.0.0 release information	Added emr-spark-8.0.0 (Apache Spark 4.0.2) release information. First GA release of Spark 4.x on EMR on EKS.	May 2026
Added EMR 7.13.0 release information	Added EMR 7.13.0 release information with Python 3.11 default for PySpark and application upgrades	April 2026
Added EMR 7.12.0 release information	Added EMR 7.12.0 release information with Iceberg Materialized Views and Hudi Full Table Access	November 2025
Added EMR 7.11.0 release information	Added EMR 7.11.0 release information with SageMaker Unified Studio integration	November 2025
Update content	Managed policies for Amazon EMR on EKS – Additional permissions for AmazonEMRContainersServiceRolePolicy .	February 3, 2025
New release	Amazon EMR on EKS 7.6.0 releases	January 10, 2025
New release	Amazon EMR on EKS 7.5.0 releases	November 21, 2024
New release	Amazon EMR on EKS 7.4.0 releases	November 13, 2024
New release	Amazon EMR on EKS 7.3.0 releases	October 16, 2024
New release	Amazon EMR on EKS 7.2.0 releases	July 25, 2024

Change	Description	Date
New release	Amazon EMR on EKS 7.1.0 releases	April 17, 2024
New release	Amazon EMR on EKS 7.0.0 releases	December 22, 2023
New release	Amazon EMR on EKS 6.15.0 releases	November 17, 2023
New release	Amazon EMR on EKS 6.14.0 releases	October 17, 2023
Update content	Rename "managed endpoints" to interactive endpoints; Interactive endpoints general availability	September 29, 2023
New release	Amazon EMR on EKS 6.13.0 releases , and public preview docs for Running Flink jobs with Amazon EMR on EKS	September 12, 2023
New release	Amazon EMR on EKS 6.12.0 releases	July 21, 2023
New content	Added Using Volcano as a custom scheduler for Apache Spark on Amazon EMR on EKS	June 13, 2023
New content	Added Using Volcano as a custom scheduler for Apache Spark on Amazon EMR on EKS	June 13, 2023
New content	Added Using Spark container log rotation	June 12, 2023
Update content	Updated the custom image documentation for finding base image information in the Amazon ECR Public Gallery.	June 8, 2023
New release	Amazon EMR on EKS 6.11.0 releases	June 8, 2023

Change	Description	Date
New content	Added Running Spark jobs with the Spark operator and re-organized the Job Runs sections under Running Spark jobs with Amazon EMR on EKS .	June 5, 2023
New content	Added two sections: Using vertical autoscaling with Amazon EMR Spark jobs and Using self-hosted Jupyter notebooks	May 4, 2023
Document history page	Created a document history page for Amazon EMR on EKS.	March 13, 2023
Managed policies page	Created a managed policies page for Amazon EMR on EKS.	March 13, 2023