
AWS Batch
API Reference
API Version 2016-08-10



AWS Batch: API Reference

Copyright © 2018 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

Welcome	1
Actions	2
CancelJob	3
Request Syntax	3
URI Request Parameters	3
Request Body	3
Response Syntax	3
Response Elements	3
Errors	4
Example	4
See Also	5
CreateComputeEnvironment	6
Request Syntax	6
URI Request Parameters	6
Request Body	6
Response Syntax	7
Response Elements	8
Errors	8
Examples	8
See Also	10
CreateJobQueue	12
Request Syntax	12
URI Request Parameters	12
Request Body	12
Response Syntax	13
Response Elements	13
Errors	13
Examples	14
See Also	15
DeleteComputeEnvironment	17
Request Syntax	17
URI Request Parameters	17
Request Body	17
Response Syntax	17
Response Elements	17
Errors	17
Example	18
See Also	18
DeleteJobQueue	20
Request Syntax	20
URI Request Parameters	20
Request Body	20
Response Syntax	20
Response Elements	20
Errors	20
Example	21
See Also	21
DeregisterJobDefinition	23
Request Syntax	23
URI Request Parameters	23
Request Body	23
Response Syntax	23
Response Elements	23
Errors	23

Example	24
See Also	24
DescribeComputeEnvironments	26
Request Syntax	26
URI Request Parameters	26
Request Body	26
Response Syntax	27
Response Elements	27
Errors	28
Example	28
See Also	29
DescribeJobDefinitions	30
Request Syntax	30
URI Request Parameters	30
Request Body	30
Response Syntax	31
Response Elements	32
Errors	32
Example	33
See Also	34
DescribeJobQueues	35
Request Syntax	35
URI Request Parameters	35
Request Body	35
Response Syntax	36
Response Elements	36
Errors	36
Example	37
See Also	38
DescribeJobs	39
Request Syntax	39
URI Request Parameters	39
Request Body	39
Response Syntax	39
Response Elements	41
Errors	41
Example	41
See Also	43
ListJobs	44
Request Syntax	44
URI Request Parameters	44
Request Body	44
Response Syntax	45
Response Elements	45
Errors	46
Examples	46
See Also	47
RegisterJobDefinition	49
Request Syntax	49
URI Request Parameters	50
Request Body	50
Response Syntax	51
Response Elements	51
Errors	51
Examples	51
See Also	53
SubmitJob	55

Request Syntax	55
URI Request Parameters	55
Request Body	55
Response Syntax	57
Response Elements	57
Errors	57
Example	58
See Also	58
TerminateJob	60
Request Syntax	60
URI Request Parameters	60
Request Body	60
Response Syntax	60
Response Elements	60
Errors	61
Example	61
See Also	62
UpdateComputeEnvironment	63
Request Syntax	63
URI Request Parameters	63
Request Body	63
Response Syntax	64
Response Elements	64
Errors	64
Example	65
See Also	65
UpdateJobQueue	67
Request Syntax	67
URI Request Parameters	67
Request Body	67
Response Syntax	68
Response Elements	68
Errors	68
Example	68
See Also	69
Data Types	70
ArrayProperties	71
Contents	71
See Also	71
ArrayPropertiesDetail	72
Contents	72
See Also	72
ArrayPropertiesSummary	73
Contents	73
See Also	73
AttemptContainerDetail	74
Contents	74
See Also	74
AttemptDetail	76
Contents	76
See Also	76
ComputeEnvironmentDetail	77
Contents	77
See Also	78
ComputeEnvironmentOrder	79
Contents	79
See Also	79

ComputeResource	80
Contents	80
See Also	81
ComputeResourceUpdate	83
Contents	83
See Also	83
ContainerDetail	84
Contents	84
See Also	86
ContainerOverrides	87
Contents	87
See Also	87
ContainerProperties	88
Contents	88
See Also	90
ContainerSummary	91
Contents	91
See Also	91
Host	92
Contents	92
See Also	92
JobDefinition	93
Contents	93
See Also	94
JobDependency	95
Contents	95
See Also	95
JobDetail	96
Contents	96
See Also	98
JobQueueDetail	99
Contents	99
See Also	100
JobSummary	101
Contents	101
See Also	102
JobTimeout	103
Contents	103
See Also	103
KeyValuePair	104
Contents	104
See Also	104
MountPoint	105
Contents	105
See Also	105
RetryStrategy	106
Contents	106
See Also	106
Ulimit	107
Contents	107
See Also	107
Volume	108
Contents	108
See Also	108
Common Parameters	109

Welcome

AWS Batch enables you to run batch computing workloads on the AWS Cloud. Batch computing is a common way for developers, scientists, and engineers to access large amounts of compute resources, and AWS Batch removes the undifferentiated heavy lifting of configuring and managing the required infrastructure. AWS Batch will be familiar to users of traditional batch computing software. This service can efficiently provision resources in response to jobs submitted in order to eliminate capacity constraints, reduce compute costs, and deliver results quickly.

As a fully managed service, AWS Batch enables developers, scientists, and engineers to run batch computing workloads of any scale. AWS Batch automatically provisions compute resources and optimizes the workload distribution based on the quantity and scale of the workloads. With AWS Batch, there is no need to install or manage batch computing software, which allows you to focus on analyzing results and solving problems. AWS Batch reduces operational complexities, saves time, and reduces costs, which makes it easy for developers, scientists, and engineers to run their batch jobs in the AWS Cloud.

This document was last published on June 16, 2018.

Actions

The following actions are supported:

- [CancelJob](#) (p. 3)
- [CreateComputeEnvironment](#) (p. 6)
- [CreateJobQueue](#) (p. 12)
- [DeleteComputeEnvironment](#) (p. 17)
- [DeleteJobQueue](#) (p. 20)
- [DeregisterJobDefinition](#) (p. 23)
- [DescribeComputeEnvironments](#) (p. 26)
- [DescribeJobDefinitions](#) (p. 30)
- [DescribeJobQueues](#) (p. 35)
- [DescribeJobs](#) (p. 39)
- [ListJobs](#) (p. 44)
- [RegisterJobDefinition](#) (p. 49)
- [SubmitJob](#) (p. 55)
- [TerminateJob](#) (p. 60)
- [UpdateComputeEnvironment](#) (p. 63)
- [UpdateJobQueue](#) (p. 67)

CancelJob

CancelJob cancels a job in an AWS Batch job queue. Jobs that are in the `SUBMITTED`, `PENDING`, or `RUNNABLE` state are cancelled. Jobs that have progressed to `STARTING` or `RUNNING` are not cancelled (but the API operation still succeeds, even if no job is cancelled); these jobs must be terminated with the [TerminateJob \(p. 60\)](#) operation.

Request Syntax

```
POST /v1/canceljob HTTP/1.1
Content-type: application/json

{
  "jobId": "string",
  "reason": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[jobId \(p. 3\)](#)

The AWS Batch job ID of the job to cancel.

Type: String

Required: Yes

[reason \(p. 3\)](#)

A message to attach to the job that explains the reason for canceling it. This message is returned by future [DescribeJobs \(p. 39\)](#) operations on the job. This message is also recorded in the AWS Batch activity logs.

Type: String

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Example

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example cancels a job with the specified job ID.

Sample Request

```
POST /v1/canceljob HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 78
Authorization: AUTHPARAMS
X-Amz-Date: 20161130T001258Z
User-Agent: aws-cli/1.11.22 Python/2.7.12 Darwin/16.1.0 botocore/1.4.79

{
  "reason": "Cancelling job.",
  "jobId": "1d828f65-7a4d-42e8-996d-3b900ed59dc4"
}
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 2
Connection: keep-alive
Date: Wed, 30 Nov 2016 00:12:59 GMT
x-amzn-RequestId: c0049e91-b691-11e6-964d-89ad0cc99f8e
X-Amzn-Trace-Id: Root=1-583e198a-cc8df0f4fac14f0d51777093
X-Cache: Miss from cloudfront
Via: 1.1 bfdd5909914586f5bc4851846228c27f.cloudfront.net (CloudFront)
X-Amz-Cf-Id: whn1dX1uTx34Lvao7-7ZdkDXEbCZ_sjn3v3hHVFgbo1ORJtXyeggSw==
```

```
{ }
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

CreateComputeEnvironment

Creates an AWS Batch compute environment. You can create `MANAGED` or `UNMANAGED` compute environments.

In a managed compute environment, AWS Batch manages the compute resources within the environment, based on the compute resources that you specify. Instances launched into a managed compute environment use a recent, approved version of the Amazon ECS-optimized AMI. You can choose to use Amazon EC2 On-Demand Instances or Spot Instances in your managed compute environment. You can optionally set a maximum price so that Spot Instances only launch when the Spot Instance price is below a specified percentage of the On-Demand price.

In an unmanaged compute environment, you can manage your own compute resources. This provides more compute resource configuration options, such as using a custom AMI, but you must ensure that your AMI meets the Amazon ECS container instance AMI specification. For more information, see [Container Instance AMIs](#) in the *Amazon Elastic Container Service Developer Guide*. After you have created your unmanaged compute environment, you can use the [DescribeComputeEnvironments \(p. 26\)](#) operation to find the Amazon ECS cluster that is associated with it and then manually launch your container instances into that Amazon ECS cluster. For more information, see [Launching an Amazon ECS Container Instance](#) in the *Amazon Elastic Container Service Developer Guide*.

Request Syntax

```
POST /v1/createcomputeenvironment HTTP/1.1  
Content-type: application/json
```

```
{  
  "computeEnvironmentName": "string",  
  "computeResources": {  
    "bidPercentage": number,  
    "desiredvCpus": number,  
    "ec2KeyPair": "string",  
    "imageId": "string",  
    "instanceRole": "string",  
    "instanceTypes": [ "string" ],  
    "maxvCpus": number,  
    "minvCpus": number,  
    "securityGroupIds": [ "string" ],  
    "spotIamFleetRole": "string",  
    "subnets": [ "string" ],  
    "tags": {  
      "string" : "string"  
    },  
    "type": "string"  
  },  
  "serviceRole": "string",  
  "state": "string",  
  "type": "string"  
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

computeEnvironmentName (p. 6)

The name for your compute environment. Up to 128 letters (uppercase and lowercase), numbers, hyphens, and underscores are allowed.

Type: String

Required: Yes

computeResources (p. 6)

Details of the compute resources managed by the compute environment. This parameter is required for managed compute environments.

Type: [ComputeResource \(p. 80\)](#) object

Required: No

serviceRole (p. 6)

The full Amazon Resource Name (ARN) of the IAM role that allows AWS Batch to make calls to other AWS services on your behalf.

If your specified role has a path other than /, then you must either specify the full role ARN (this is recommended) or prefix the role name with the path.

Note

Depending on how you created your AWS Batch service role, its ARN may contain the `service-role` path prefix. When you only specify the name of the service role, AWS Batch assumes that your ARN does not use the `service-role` path prefix. Because of this, we recommend that you specify the full ARN of your service role when you create compute environments.

Type: String

Required: Yes

state (p. 6)

The state of the compute environment. If the state is `ENABLED`, then the compute environment accepts jobs from a queue and can scale out automatically based on queues.

Type: String

Valid Values: `ENABLED` | `DISABLED`

Required: No

type (p. 6)

The type of the compute environment.

Type: String

Valid Values: `MANAGED` | `UNMANAGED`

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "computeEnvironmentArn": "string",
```

```
"computeEnvironmentName": "string"  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

computeEnvironmentArn (p. 7)

The Amazon Resource Name (ARN) of the compute environment.

Type: String

computeEnvironmentName (p. 7)

The name of the compute environment.

Type: String

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Examples

In the following example or examples, the Authorization header contents (`AUTHPARAMS`) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example creates a managed compute environment with specific C4 instance types that are launched on demand. The compute environment is called `C4OnDemand`.

Sample Request

```
POST /v1/createcomputeenvironment HTTP/1.1  
Host: batch.us-east-1.amazonaws.com  
Accept-Encoding: identity
```

```
Content-Length: 538
Authorization: AUTHPARAMS
X-Amz-Date: 20161128T223128Z
User-Agent: aws-cli/1.11.21 Python/2.7.12 Darwin/16.1.0 botocore/1.4.78

{
  "computeEnvironmentName": "C4OnDemand",
  "state": "ENABLED",
  "type": "MANAGED",
  "computeResources": {
    "subnets": [
      "subnet-220c0e0a",
      "subnet-1a95556d",
      "subnet-978f6dce"
    ],
    "tags": {
      "Name": "Batch Instance - C4OnDemand"
    },
    "desiredvCpus": 48,
    "minvCpus": 0,
    "instanceTypes": [
      "c4.large",
      "c4.xlarge",
      "c4.2xlarge",
      "c4.4xlarge",
      "c4.8xlarge"
    ],
    "securityGroupIds": [
      "sg-cf5093b2"
    ],
    "instanceRole": "ecsInstanceRole",
    "maxvCpus": 128,
    "type": "EC2",
    "ec2KeyPair": "id_rsa"
  },
  "serviceRole": "arn:aws:iam::012345678910:role/AWSBatchServiceRole"
}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 28 Nov 2016 22:31:28 GMT
Content-Type: application/json
Content-Length: 133
Connection: keep-alive
x-amzn-RequestId: 67558123-b5ba-11e6-b909-370261104d69
X-Amzn-Trace-Id: Root=1-583cb040-93d52243313314b33c0c1e97
X-Cache: Miss from cloudfront
Via: 1.1 7e587c722adb25336835ccb4e5814e4e.cloudfront.net (CloudFront)
X-Amz-Cf-Id: GwQRsxvmiuj1HYwbYq9MAEsQfJpN6BknGQlNX1jAd5qLQFXyHBwOUQ==

{
  "computeEnvironmentName": "C4OnDemand",
  "computeEnvironmentArn": "arn:aws:batch:us-east-1:012345678910:compute-environment/C4OnDemand"
}
```

Example

This example creates a managed compute environment with the M4 instance type that is launched when the Spot Instance price is at or below 20% of the On-Demand price for the instance type. The compute environment is called M4Spot.

Sample Request

```
POST /v1/createcomputeenvironment HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 568
Authorization: AUTHPARAMS
X-Amz-Date: 20161128T223813Z
User-Agent: aws-cli/1.11.21 Python/2.7.12 Darwin/16.1.0 botocore/1.4.78

{
  "computeEnvironmentName": "M4Spot",
  "state": "ENABLED",
  "type": "MANAGED",
  "computeResources": {
    "subnets": [
      "subnet-220c0e0a",
      "subnet-1a95556d",
      "subnet-978f6dce"
    ],
    "type": "SPOT",
    "spotIamFleetRole": "arn:aws:iam::012345678910:role/aws-ec2-spot-fleet-role",
    "tags": {
      "Name": "Batch Instance - M4Spot"
    },
    "desiredvCpus": 4,
    "minvCpus": 0,
    "instanceTypes": [
      "m4"
    ],
    "securityGroupIds": [
      "sg-cf5093b2"
    ],
    "instanceRole": "ecsInstanceRole",
    "maxvCpus": 128,
    "bidPercentage": 20,
    "ec2KeyPair": "id_rsa"
  },
  "serviceRole": "arn:aws:iam::012345678910:role/AWSBatchServiceRole"
}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 28 Nov 2016 22:38:16 GMT
Content-Type: application/json
Content-Length: 125
Connection: keep-alive
x-amzn-RequestId: 59422e20-b5bb-11e6-9a64-53057c7adce9
X-Amzn-Trace-Id: Root=1-583cb1d6-b71bcabf4f10bae0f3ade63b
X-Cache: Miss from cloudfront
Via: 1.1 8455edd9286a1292a39c993fdeccce65.cloudfront.net (CloudFront)
X-Amz-Cf-Id: 4mklLyUpygUko86fMNzPgA8_D64lSwPmG6iIKhAZkGpOp2e-3cKg_w==

{
  "computeEnvironmentName": "M4Spot",
  "computeEnvironmentArn": "arn:aws:batch:us-east-1:012345678910:compute-environment/M4Spot"
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

CreateJobQueue

Creates an AWS Batch job queue. When you create a job queue, you associate one or more compute environments to the queue and assign an order of preference for the compute environments.

You also set a priority to the job queue that determines the order in which the AWS Batch scheduler places jobs onto its associated compute environments. For example, if a compute environment is associated with more than one job queue, the job queue with a higher priority is given preference for scheduling jobs to that compute environment.

Request Syntax

```
POST /v1/createjobqueue HTTP/1.1
Content-type: application/json

{
  "computeEnvironmentOrder": [
    {
      "computeEnvironment": "string",
      "order": number
    }
  ],
  "jobQueueName": "string",
  "priority": number,
  "state": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

computeEnvironmentOrder (p. 12)

The set of compute environments mapped to a job queue and their order relative to each other. The job scheduler uses this parameter to determine which compute environment should execute a given job. Compute environments must be in the `VALID` state before you can associate them with a job queue. You can associate up to three compute environments with a job queue.

Type: Array of [ComputeEnvironmentOrder](#) (p. 79) objects

Required: Yes

jobQueueName (p. 12)

The name of the job queue.

Type: String

Required: Yes

priority (p. 12)

The priority of the job queue. Job queues with a higher priority (or a higher integer value for the `priority` parameter) are evaluated first when associated with same compute environment. Priority

is determined in descending order, for example, a job queue with a priority value of 10 is given scheduling preference over a job queue with a priority value of 1.

Type: Integer

Required: Yes

[state \(p. 12\)](#)

The state of the job queue. If the job queue state is `ENABLED`, it is able to accept jobs.

Type: String

Valid Values: `ENABLED` | `DISABLED`

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "jobQueueArn": "string",
  "jobQueueName": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[jobQueueArn \(p. 13\)](#)

The Amazon Resource Name (ARN) of the job queue.

Type: String

[jobQueueName \(p. 13\)](#)

The name of the job queue.

Type: String

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Examples

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example creates a job queue called `LowPriority` that uses the `M4Spot` compute environment.

Sample Request

```
POST /v1/createjobqueue HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 142
Authorization: AUTHPARAMS
X-Amz-Date: 20161128T234201Z
User-Agent: aws-cli/1.11.21 Python/2.7.12 Darwin/16.1.0 botocore/1.4.78

{
  "priority": 1,
  "state": "ENABLED",
  "computeEnvironmentOrder": [
    {
      "computeEnvironment": "M4Spot",
      "order": 1
    }
  ],
  "jobQueueName": "LowPriority"
}
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 105
Connection: keep-alive
Date: Mon, 28 Nov 2016 23:42:02 GMT
x-amzn-RequestId: 42a92eee-b5c4-11e6-b15b-8167bcebef9d
X-Amzn-Trace-Id: Root=1-583cc0c9-f47ba831b55e794291697c39
X-Cache: Miss from cloudfront
Via: 1.1 a44b4468444ef3ee67472bd5c5016098.cloudfront.net (CloudFront)
X-Amz-Cf-Id: bz9IuCM5FNkDfge5y-Zw7nFEjDdTHDYFwbEY2AKUqrt9l2XeKUcuyA==

{
  "jobQueueName": "LowPriority",
  "jobQueueArn": "arn:aws:batch:us-east-1:012345678910:job-queue/LowPriority"
}
```

Example

This example creates a job queue called `HighPriority` that uses the `C4OnDemand` compute environment with an order of 1 and the `M4Spot` compute environment with an order of 2.

Sample Request

```
POST /v1/createjobqueue HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 192
Authorization: AUTHPARAMS
X-Amz-Date: 20161128T234933Z
User-Agent: aws-cli/1.11.21 Python/2.7.12 Darwin/16.1.0 botocore/1.4.78

{
  "priority": 10,
  "state": "ENABLED",
  "computeEnvironmentOrder": [
    {
      "computeEnvironment": "C4OnDemand",
      "order": 1
    },
    {
      "computeEnvironment": "M4Spot",
      "order": 2
    }
  ],
  "jobQueueName": "HighPriority"
}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 28 Nov 2016 23:49:34 GMT
Content-Type: application/json
Content-Length: 107
Connection: keep-alive
x-amzn-RequestId: 503cdf30-b5c5-11e6-b7a0-a325cf93a1c0
X-Amzn-Trace-Id: Root=1-583cc28e-ca43f4d9fb8dc73ebb99dd67
X-Cache: Miss from cloudfront
Via: 1.1 e81bbc86832b655de5b9a19317ad01.cloudfront.net (CloudFront)
X-Amz-Cf-Id: 8NB20odDPMaKy9zHa6GPaGN_r562QsynDTRYPUhKwHSvQrMG70IHSQ==

{
  "jobQueueName": "HighPriority",
  "jobQueueArn": "arn:aws:batch:us-east-1:012345678910:job-queue/HighPriority"
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DeleteComputeEnvironment

Deletes an AWS Batch compute environment.

Before you can delete a compute environment, you must set its state to `DISABLED` with the [UpdateComputeEnvironment \(p. 63\)](#) API operation and disassociate it from any job queues with the [UpdateJobQueue \(p. 67\)](#) API operation.

Request Syntax

```
POST /v1/deletecomputeenvironment HTTP/1.1
Content-type: application/json

{
  "computeEnvironment": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

computeEnvironment (p. 17)

The name or Amazon Resource Name (ARN) of the compute environment to delete.

Type: String

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Example

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example deletes the P3OnDemand compute environment.

Sample Request

```
POST /v1/deletecomputeenvironment HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 36
Authorization: AUTHPARAMS
X-Amz-Date: 20161128T202219Z
User-Agent: aws-cli/1.11.21 Python/2.7.12 Darwin/16.1.0 botocore/1.4.78

{
  "computeEnvironment": "P3OnDemand"
}
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 2
Connection: keep-alive
Date: Mon, 28 Nov 2016 20:22:20 GMT
x-amzn-RequestId: 5cffa6c9-b5a8-11e6-a551-27cf529560ed
X-Amzn-Trace-Id: Root=1-583c91fc-e3864c0561e747945eca7135
X-Cache: Miss from cloudfront
Via: 1.1 b63769e2d89c89274acd908e4bfc9f4.cloudfront.net (CloudFront)
X-Amz-Cf-Id: mqHP9krdcbsbt0pivub4bJEMO_XCTTfENz0xPwwye-USu1CVGlj-nw==

{}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)

- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DeleteJobQueue

Deletes the specified job queue. You must first disable submissions for a queue with the [UpdateJobQueue \(p. 67\)](#) operation. All jobs in the queue are terminated when you delete a job queue.

It is not necessary to disassociate compute environments from a queue before submitting a DeleteJobQueue request.

Request Syntax

```
POST /v1/deletejobqueue HTTP/1.1
Content-type: application/json

{
  "jobQueue": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[jobQueue \(p. 20\)](#)

The short name or full Amazon Resource Name (ARN) of the queue to delete.

Type: String

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Example

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example deletes the GPGPU job queue.

Sample Request

```
POST /v1/deletejobqueue HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 21
Authorization: AUTHPARAMS
X-Amz-Date: 20161128T201857Z
User-Agent: aws-cli/1.11.21 Python/2.7.12 Darwin/16.1.0 botocore/1.4.78

{
  "jobQueue": "GPGPU"
}
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 2
Connection: keep-alive
Date: Mon, 28 Nov 2016 20:18:57 GMT
x-amzn-RequestId: e44df9d5-b5a7-11e6-a551-27cf529560ed
X-Amzn-Trace-Id: Root=1-583c9131-d8de3fb3afbale0ae32e9307
X-Cache: Miss from cloudfront
Via: 1.1 56908f89e8d17ba579c0607313114955.cloudfront.net (CloudFront)
X-Amz-Cf-Id: UnpbX7PjdrV3N-Y79pD6eV3DfqYUXdEx3HAI9VYhUZ8h7yRbi5_ZVQ==

{}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)

- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DeregisterJobDefinition

Deregisters an AWS Batch job definition.

Request Syntax

```
POST /v1/deregisterjobdefinition HTTP/1.1
Content-type: application/json

{
  "jobDefinition": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

jobDefinition (p. 23)

The name and revision (`name:revision`) or full Amazon Resource Name (ARN) of the job definition to deregister.

Type: String

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Example

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example deregisters a job definition called `sleep10`.

Sample Request

```
POST /v1/deregisterjobdefinition HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 28
Authorization: AUTHPARAMS
X-Amz-Date: 20161128T215745Z
User-Agent: aws-cli/1.11.21 Python/2.7.12 Darwin/16.1.0 botocore/1.4.78

{
  "jobDefinition": "sleep10"
}
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 2
Connection: keep-alive
Date: Mon, 28 Nov 2016 21:57:45 GMT
x-amzn-RequestId: b18944e7-b5b5-11e6-969f-6b235d2d633d
X-Amzn-Trace-Id: Root=1-583ca859-ac12b61650100be5a6f0af35
X-Cache: Miss from cloudfront
Via: 1.1 e892630891779ff1ccadccf205a776f3.cloudfront.net (CloudFront)
X-Amz-Cf-Id: wKAY_NOTbvY8PFcmo1aGja0xqGLxsTJgEtuc1KosPYAPYL8icYwvKw==

{}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DescribeComputeEnvironments

Describes one or more of your compute environments.

If you are using an unmanaged compute environment, you can use the `DescribeComputeEnvironment` operation to determine the `ecsClusterArn` that you should launch your Amazon ECS container instances into.

Request Syntax

```
POST /v1/describecomputeenvironments HTTP/1.1
Content-type: application/json

{
  "computeEnvironments": [ "string" ],
  "maxResults": number,
  "nextToken": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

`computeEnvironments` (p. 26)

A list of up to 100 compute environment names or full Amazon Resource Name (ARN) entries.

Type: Array of strings

Required: No

`maxResults` (p. 26)

The maximum number of cluster results returned by `DescribeComputeEnvironments` in paginated output. When this parameter is used, `DescribeComputeEnvironments` only returns `maxResults` results in a single page along with a `nextToken` response element. The remaining results of the initial request can be seen by sending another `DescribeComputeEnvironments` request with the returned `nextToken` value. This value can be between 1 and 100. If this parameter is not used, then `DescribeComputeEnvironments` returns up to 100 results and a `nextToken` value if applicable.

Type: Integer

Required: No

`nextToken` (p. 26)

The `nextToken` value returned from a previous paginated `DescribeComputeEnvironments` request where `maxResults` was used and the results exceeded the value of that parameter. Pagination continues from the end of the previous results that returned the `nextToken` value. This value is `null` when there are no more results to return.

Note

This token should be treated as an opaque identifier that is only used to retrieve the next items in a list and not for other programmatic purposes.

Type: String

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "computeEnvironments": [
    {
      "computeEnvironmentArn": "string",
      "computeEnvironmentName": "string",
      "computeResources": {
        "bidPercentage": number,
        "desiredvCpus": number,
        "ec2KeyPair": "string",
        "imageId": "string",
        "instanceRole": "string",
        "instanceTypes": [ "string" ],
        "maxvCpus": number,
        "minvCpus": number,
        "securityGroupIds": [ "string" ],
        "spotIamFleetRole": "string",
        "subnets": [ "string" ],
        "tags": {
          "string" : "string"
        },
        "type": "string"
      },
      "ecsClusterArn": "string",
      "serviceRole": "string",
      "state": "string",
      "status": "string",
      "statusReason": "string",
      "type": "string"
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

computeEnvironments (p. 27)

The list of compute environments.

Type: Array of [ComputeEnvironmentDetail](#) (p. 77) objects

nextToken (p. 27)

The nextToken value to include in a future DescribeComputeEnvironments request. When the results of a DescribeJobDefinitions request exceed maxResults, this value can be used to retrieve the next page of results. This value is null when there are no more results to return.

Type: String

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Example

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example describes the P3OnDemand compute environment.

Sample Request

```
POST /v1/describecomputeenvironments HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 39
Authorization: AUTHPARAMS
X-Amz-Date: 20161128T193355Z
User-Agent: aws-cli/1.11.21 Python/2.7.12 Darwin/16.1.0 botocore/1.4.78

{
  "computeEnvironments": [
    "P3OnDemand"
  ]
}
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 742
Connection: keep-alive
Date: Mon, 28 Nov 2016 19:33:56 GMT
x-amzn-RequestId: 9a3315bc-b5a1-11e6-b942-8d20a4fbbd5a
X-Amzn-Trace-Id: Root=1-583c86a4-abecba6b67c6a740c8891289
X-Cache: Miss from cloudfront
Via: 1.1 56908f89e8d17ba579c0607313114955.cloudfront.net (CloudFront)
```

```
X-Amz-Cf-Id: FbgslaatWhj_yGhfkSCTPpPtjjuVuFOBns-kK5EsaasYQC5p2FnpiQ==

{
  "computeEnvironments": [{
    "computeEnvironmentName": "P3OnDemand",
    "computeEnvironmentArn": "arn:aws:batch:us-east-1:012345678910:compute-environment/
P3OnDemand",
    "ecsClusterArn": "arn:aws:ecs:us-east-1:012345678910:cluster/P3OnDemand_Batch_2c06f29d-
d1fe-3a49-879d-42394c86effc",
    "type": "MANAGED",
    "state": "ENABLED",
    "status": "VALID",
    "statusReason": "ComputeEnvironment Healthy",
    "computeResources": {
      "type": "EC2",
      "minvCpus": 0,
      "maxvCpus": 128,
      "desiredvCpus": 48,
      "instanceTypes": ["p3"],
      "subnets": ["subnet-220c0e0a", "subnet-1a95556d", "subnet-978f6dce"],
      "securityGroupIds": ["sg-cf5093b2"],
      "ec2KeyPair": "id_rsa",
      "instanceRole": "ecsInstanceRole",
      "tags": {
        "Name": "Batch Instance - P3OnDemand"
      }
    },
    "serviceRole": "arn:aws:iam::012345678910:role/AWSBatchServiceRole"
  }]
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DescribeJobDefinitions

Describes a list of job definitions. You can specify a `status` (such as `ACTIVE`) to only return job definitions that match that status.

Request Syntax

```
POST /v1/describejobdefinitions HTTP/1.1
Content-type: application/json

{
  "jobDefinitionName": "string",
  "jobDefinitions": [ "string" ],
  "maxResults": number,
  "nextToken": "string",
  "status": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

jobDefinitionName (p. 30)

The name of the job definition to describe.

Type: String

Required: No

jobDefinitions (p. 30)

A space-separated list of up to 100 job definition names or full Amazon Resource Name (ARN) entries.

Type: Array of strings

Required: No

maxResults (p. 30)

The maximum number of results returned by `DescribeJobDefinitions` in paginated output. When this parameter is used, `DescribeJobDefinitions` only returns `maxResults` results in a single page along with a `nextToken` response element. The remaining results of the initial request can be seen by sending another `DescribeJobDefinitions` request with the returned `nextToken` value. This value can be between 1 and 100. If this parameter is not used, then `DescribeJobDefinitions` returns up to 100 results and a `nextToken` value if applicable.

Type: Integer

Required: No

[nextToken \(p. 30\)](#)

The `nextToken` value returned from a previous paginated `DescribeJobDefinitions` request where `maxResults` was used and the results exceeded the value of that parameter. Pagination continues from the end of the previous results that returned the `nextToken` value. This value is `null` when there are no more results to return.

Note

This token should be treated as an opaque identifier that is only used to retrieve the next items in a list and not for other programmatic purposes.

Type: String

Required: No

[status \(p. 30\)](#)

The status with which to filter job definitions.

Type: String

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "jobDefinitions": [
    {
      "containerProperties": {
        "command": [ "string" ],
        "environment": [
          {
            "name": "string",
            "value": "string"
          }
        ],
        "image": "string",
        "jobRoleArn": "string",
        "memory": number,
        "mountPoints": [
          {
            "containerPath": "string",
            "readOnly": boolean,
            "sourceVolume": "string"
          }
        ],
        "privileged": boolean,
        "readonlyRootFilesystem": boolean,
        "ulimits": [
          {
            "hardLimit": number,
            "name": "string",
            "softLimit": number
          }
        ],
        "user": "string",
        "vcpus": number,
        "volumes": [
          {
```

```

        "host": {
            "sourcePath": "string"
        },
        "name": "string"
    }
]
},
"jobDefinitionArn": "string",
"jobDefinitionName": "string",
"parameters": {
    "string": "string"
},
"retryStrategy": {
    "attempts": number
},
"revision": number,
"status": "string",
"timeout": {
    "attemptDurationSeconds": number
},
"type": "string"
}
],
"nextToken": "string"
}

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[jobDefinitions \(p. 31\)](#)

The list of job definitions.

Type: Array of [JobDefinition \(p. 93\)](#) objects

[nextToken \(p. 31\)](#)

The `nextToken` value to include in a future `DescribeJobDefinitions` request. When the results of a `DescribeJobDefinitions` request exceed `maxResults`, this value can be used to retrieve the next page of results. This value is `null` when there are no more results to return.

Type: String

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Example

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example describes all of your active job definitions.

Sample Request

```
POST /v1/describejobdefinitions HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 20
Authorization: AUTHPARAMS
X-Amz-Date: 20161128T221855Z
User-Agent: aws-cli/1.11.21 Python/2.7.12 Darwin/16.1.0 botocore/1.4.78

{
  "status": "ACTIVE"
}
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 351
Connection: keep-alive
Date: Mon, 28 Nov 2016 22:18:55 GMT
x-amzn-RequestId: a6856a57-b5b8-11e6-b057-0594e7a2cf06
X-Amzn-Trace-Id: Root=1-583cad4f-d67ae20374b79d9adbaca1fb
X-Cache: Miss from cloudfront
Via: 1.1 688936cc730f240888e6a59a81892a3d.cloudfront.net (CloudFront)
X-Amz-Cf-Id: hd-CAMqfaCJt-1oH7tBu9j5c-IhLQuMjFHFpck6F0MMt5CBea8mQBQ==

{
  "jobDefinitions": [{
    "jobDefinitionName": "sleep60",
    "jobDefinitionArn": "arn:aws:batch:us-east-1:012345678910:job-definition/sleep60:1",
    "revision": 1,
    "status": "ACTIVE",
    "type": "container",
    "containerProperties": {
      "image": "busybox",
      "vcpus": 1,
      "memory": 128,
      "command": ["sleep", "60"],
      "volumes": [],
      "environment": [],
      "mountPoints": [],
      "ulimits": []
    }
  }]
}
```

}

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DescribeJobQueues

Describes one or more of your job queues.

Request Syntax

```
POST /v1/describejobqueues HTTP/1.1
Content-type: application/json

{
  "jobQueues": [ "string" ],
  "maxResults": number,
  "nextToken": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[jobQueues \(p. 35\)](#)

A list of up to 100 queue names or full queue Amazon Resource Name (ARN) entries.

Type: Array of strings

Required: No

[maxResults \(p. 35\)](#)

The maximum number of results returned by `DescribeJobQueues` in paginated output. When this parameter is used, `DescribeJobQueues` only returns `maxResults` results in a single page along with a `nextToken` response element. The remaining results of the initial request can be seen by sending another `DescribeJobQueues` request with the returned `nextToken` value. This value can be between 1 and 100. If this parameter is not used, then `DescribeJobQueues` returns up to 100 results and a `nextToken` value if applicable.

Type: Integer

Required: No

[nextToken \(p. 35\)](#)

The `nextToken` value returned from a previous paginated `DescribeJobQueues` request where `maxResults` was used and the results exceeded the value of that parameter. Pagination continues from the end of the previous results that returned the `nextToken` value. This value is `null` when there are no more results to return.

Note

This token should be treated as an opaque identifier that is only used to retrieve the next items in a list and not for other programmatic purposes.

Type: String

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "jobQueues": [
    {
      "computeEnvironmentOrder": [
        {
          "computeEnvironment": "string",
          "order": number
        }
      ],
      "jobQueueArn": "string",
      "jobQueueName": "string",
      "priority": number,
      "state": "string",
      "status": "string",
      "statusReason": "string"
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

jobQueues (p. 36)

The list of job queues.

Type: Array of [JobQueueDetail](#) (p. 99) objects

nextToken (p. 36)

The `nextToken` value to include in a future `DescribeJobQueues` request. When the results of a `DescribeJobQueues` request exceed `maxResults`, this value can be used to retrieve the next page of results. This value is `null` when there are no more results to return.

Type: String

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Example

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example describes the HighPriority job queue.

Sample Request

```
POST /v1/describejobqueues HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 31
Authorization: AUTHPARAMS
X-Amz-Date: 20161128T194731Z
User-Agent: aws-cli/1.11.21 Python/2.7.12 Darwin/16.1.0 botocore/1.4.78

{
  "jobQueues": [
    "HighPriority"
  ]
}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 28 Nov 2016 19:47:32 GMT
Content-Type: application/json
Content-Length: 336
Connection: keep-alive
x-amzn-RequestId: 8073d3c6-b5a3-11e6-848a-3577abb4ef72
X-Amzn-Trace-Id: Root=1-583c89d4-dd8e9009bdc002a2f5f6f685
X-Cache: Miss from cloudfront
Via: 1.1 7bfcc2251021d9dc94a87ff179d69731.cloudfront.net (CloudFront)
X-Amz-Cf-Id: dwf7P2pnEYCxN1C3XdApqDtqzLfjpwAjbHvskd9oUqz40Un9pvtx3Q==

{
  "jobQueues": [{
    "jobQueueName": "HighPriority",
    "jobQueueArn": "arn:aws:batch:us-east-1:012345678910:job-queue/HighPriority",
    "state": "ENABLED",
    "status": "VALID",
    "statusReason": "JobQueue Healthy",
    "priority": 10,
    "computeEnvironmentOrder": [{
      "order": 1,
      "computeEnvironment": "arn:aws:batch:us-east-1:012345678910:compute-environment/
C4OnDemand"
    }]
  }]
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DescribeJobs

Describes a list of AWS Batch jobs.

Request Syntax

```
POST /v1/describejobs HTTP/1.1
Content-type: application/json

{
  "jobs": [ "string" ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[jobs \(p. 39\)](#)

A space-separated list of up to 100 job IDs.

Type: Array of strings

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "jobs": [
    {
      "arrayProperties": {
        "index": number,
        "size": number,
        "statusSummary": {
          "string": number
        }
      },
      "attempts": [
        {
          "container": {
            "containerInstanceArn": "string",
            "exitCode": number,
            "logStreamName": "string",
            "reason": "string",
            "taskArn": "string"
          },
          "startedAt": number,
          "statusReason": "string",
          "stoppedAt": number
        }
      ]
    }
  ]
}
```

```

    }
  ],
  "container": {
    "command": [ "string" ],
    "containerInstanceArn": "string",
    "environment": [
      {
        "name": "string",
        "value": "string"
      }
    ],
    "exitCode": number,
    "image": "string",
    "jobRoleArn": "string",
    "logStreamName": "string",
    "memory": number,
    "mountPoints": [
      {
        "containerPath": "string",
        "readOnly": boolean,
        "sourceVolume": "string"
      }
    ],
    "privileged": boolean,
    "readonlyRootFilesystem": boolean,
    "reason": "string",
    "taskArn": "string",
    "ulimits": [
      {
        "hardLimit": number,
        "name": "string",
        "softLimit": number
      }
    ],
    "user": "string",
    "vcpus": number,
    "volumes": [
      {
        "host": {
          "sourcePath": "string"
        },
        "name": "string"
      }
    ]
  },
  "createdAt": number,
  "dependsOn": [
    {
      "jobId": "string",
      "type": "string"
    }
  ],
  "jobDefinition": "string",
  "jobId": "string",
  "jobName": "string",
  "jobQueue": "string",
  "parameters": {
    "string": "string"
  },
  "retryStrategy": {
    "attempts": number
  },
  "startedAt": number,
  "status": "string",
  "statusReason": "string",
  "stoppedAt": number,

```

```
        "timeout": {  
            "attemptDurationSeconds": number  
        }  
    }  
]  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[jobs \(p. 39\)](#)

The list of jobs.

Type: Array of [JobDetail \(p. 96\)](#) objects

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Example

In the following example or examples, the Authorization header contents (`AUTHPARAMS`) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example describes a job with the specified job ID.

Sample Request

```
POST /v1/describejobs HTTP/1.1  
Host: batch.us-east-1.amazonaws.com  
Accept-Encoding: identity  
Content-Length: 50
```

```
Authorization: AUTHPARAMS
X-Amz-Date: 20170327T151323Z
User-Agent: aws-cli/1.11.22 Python/2.7.12 Darwin/16.1.0 botocore/1.4.79
```

```
{
  "jobs": [
    "0668da57-1bcc-478b-bc14-5d4f1c1cef48"
  ]
}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 27 Mar 2017 15:13:13 GMT
Content-Type: application/json
Content-Length: 1883
Connection: keep-alive
x-amzn-RequestId: e5628975-12ff-11e7-ab46-a583c88f0f07
X-Amzn-Trace-Id: Root=1-58d92c09-19ef6bec3f9f5a392d25738f
X-Cache: Miss from cloudfront
Via: 1.1 8a78b675adb2cce925860f2fe4383e71.cloudfront.net (CloudFront)
X-Amz-Cf-Id: TaW8k7yrDyXHgEU2udEEOAbliIY1iPmQr4LpN8OULdqyVGR6qP0q4Q==

{
  "jobs": [
    {
      "jobName": "EchoAttemptNumber",
      "jobId": "0668da57-1bcc-478b-bc14-5d4f1c1cef48",
      "jobQueue": "arn:aws:batch:us-east-1:012345678910:job-queue/HighPriority",
      "status": "FAILED",
      "attempts": [
        {
          "container": {
            "containerInstanceArn": "arn:aws:ecs:us-east-1:012345678910:container-
instance/90bfe527-119c-494a-b8fe-f5999c66d214",
            "taskArn": "arn:aws:ecs:us-east-1:012345678910:task/af37d830-6978-4a2b-b796-
e890e9b477b3",
            "exitCode": 1
          },
          "startedAt": 1490627002951,
          "stoppedAt": 1490627003065,
          "statusReason": "Essential container in task exited"
        },
        {
          "container": {
            "containerInstanceArn": "arn:aws:ecs:us-east-1:012345678910:container-
instance/90bfe527-119c-494a-b8fe-f5999c66d214",
            "taskArn": "arn:aws:ecs:us-east-1:012345678910:task/3dfd4d0e-
a177-4798-9c13-21b7148217bc",
            "exitCode": 2
          },
          "startedAt": 1490627019948,
          "stoppedAt": 1490627020059,
          "statusReason": "Essential container in task exited"
        },
        {
          "container": {
            "containerInstanceArn": "arn:aws:ecs:us-east-1:012345678910:container-
instance/90bfe527-119c-494a-b8fe-f5999c66d214",
            "taskArn": "arn:aws:ecs:us-east-1:012345678910:task/22857040-182c-4af3-85f5-
bb2c71edd282",
            "exitCode": 3
          },
          "startedAt": 1490627034798,

```



```
        "stoppedAt": 1490627034949,
        "statusReason": "Essential container in task exited"
    }
],
"statusReason": "Essential container in task exited",
"createdAt": 1490626709525,
"retryStrategy": {
    "attempts": 3
},
"startedAt": 1490627034798,
"stoppedAt": 1490627034949,
"dependsOn": [],
"jobDefinition": "arn:aws:batch:us-east-1:012345678910:job-definition/
EchoAttemptNumber:1",
"parameters": {},
"container": {
    "image": "amazonlinux",
    "vcpus": 1,
    "memory": 2,
    "command": [
        "/bin/bash",
        "-c",
        "exit $AWS_BATCH_JOB_ATTEMPT"
    ],
    "volumes": [],
    "environment": [],
    "mountPoints": [],
    "ulimits": [],
    "exitCode": 3,
    "containerInstanceArn": "arn:aws:ecs:us-east-1:012345678910:container-
instance/90bfe527-119c-494a-b8fe-f5999c66d214",
    "taskArn": "arn:aws:ecs:us-east-1:012345678910:task/22857040-182c-4af3-85f5-
bb2c71edd282"
}
}
]
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

ListJobs

Returns a list of task jobs for a specified job queue. You can filter the results by job status with the `jobStatus` parameter. If you do not specify a status, only `RUNNING` jobs are returned.

Request Syntax

```
POST /v1/listjobs HTTP/1.1
Content-type: application/json
```

```
{
  "arrayJobId": "string",
  "jobQueue": "string",
  "jobStatus": "string",
  "maxResults": number,
  "nextToken": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[arrayJobId \(p. 44\)](#)

The job ID for an array job. Specifying an array job ID with this parameter lists all child jobs from within the specified array.

Type: String

Required: No

[jobQueue \(p. 44\)](#)

The name or full Amazon Resource Name (ARN) of the job queue with which to list jobs.

Type: String

Required: No

[jobStatus \(p. 44\)](#)

The job status with which to filter jobs in the specified queue. If you do not specify a status, only `RUNNING` jobs are returned.

Type: String

Valid Values: `SUBMITTED` | `PENDING` | `RUNNABLE` | `STARTING` | `RUNNING` | `SUCCEEDED` | `FAILED`

Required: No

[maxResults \(p. 44\)](#)

The maximum number of results returned by `ListJobs` in paginated output. When this parameter is used, `ListJobs` only returns `maxResults` results in a single page along with a `nextToken` response element. The remaining results of the initial request can be seen by sending another

`ListJobs` request with the returned `nextToken` value. This value can be between 1 and 100. If this parameter is not used, then `ListJobs` returns up to 100 results and a `nextToken` value if applicable.

Type: Integer

Required: No

[nextToken \(p. 44\)](#)

The `nextToken` value returned from a previous paginated `ListJobs` request where `maxResults` was used and the results exceeded the value of that parameter. Pagination continues from the end of the previous results that returned the `nextToken` value. This value is `null` when there are no more results to return.

Note

This token should be treated as an opaque identifier that is only used to retrieve the next items in a list and not for other programmatic purposes.

Type: String

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "jobSummaryList": [
    {
      "arrayProperties": {
        "index": number,
        "size": number
      },
      "container": {
        "exitCode": number,
        "reason": "string"
      },
      "createdAt": number,
      "jobId": "string",
      "jobName": "string",
      "startedAt": number,
      "status": "string",
      "statusReason": "string",
      "stoppedAt": number
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[jobSummaryList \(p. 45\)](#)

A list of job summaries that match the request.

Type: Array of [JobSummary](#) (p. 101) objects
[nextToken](#) (p. 45)

The `nextToken` value to include in a future `ListJobs` request. When the results of a `ListJobs` request exceed `maxResults`, this value can be used to retrieve the next page of results. This value is `null` when there are no more results to return.

Type: String

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Examples

In the following example or examples, the Authorization header contents (`AUTHPARAMS`) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example lists the running jobs in the `HighPriority` job queue.

Sample Request

```
POST /v1/listjobs HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 28
Authorization: AUTHPARAMS
X-Amz-Date: 20161129T201622Z
User-Agent: aws-cli/1.11.22 Python/2.7.12 Darwin/16.1.0 botocore/1.4.79

{
  "jobQueue": "HighPriority"
}
```

Sample Response

```
HTTP/1.1 200 OK
```

```
Content-Type: application/json
Content-Length: 89
Connection: keep-alive
Date: Tue, 29 Nov 2016 20:16:22 GMT
x-amzn-RequestId: b264511f-b670-11e6-b8df-d7216c30de5e
X-Amzn-Trace-Id: Root=1-583de216-c8d6ac58e2bd21ac2d8284cf
X-Cache: Miss from cloudfront
Via: 1.1 7f3f42df8af148df1f9f1ee7175987ad.cloudfront.net (CloudFront)
X-Amz-Cf-Id: idKR5mD8f7Luom03P9DV1bFGXsq_SIFNy_nMrTCOqZrRc0nXgHqZfg==

{
  "jobSummaryList": [{
    "jobId": "e66ff5fd-alf-4640-b1a2-0b0a142f49bb",
    "jobName": "example"
  }]
}
```

Example

This example lists jobs in the HighPriority job queue that are in the SUBMITTED job status.

Sample Request

```
POST /v1/listjobs HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 54
Authorization: AUTHPARAMS
X-Amz-Date: 20161129T201642Z
User-Agent: aws-cli/1.11.22 Python/2.7.12 Darwin/16.1.0 botocore/1.4.79

{
  "jobQueue": "HighPriority",
  "jobStatus": "SUBMITTED"
}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Tue, 29 Nov 2016 20:16:42 GMT
Content-Type: application/json
Content-Length: 89
Connection: keep-alive
x-amzn-RequestId: be15ca04-b670-11e6-aa0e-ef9532a24bfe
X-Amzn-Trace-Id: Root=1-583de22a-fdc493168642bc6d4de19ba7
X-Cache: Miss from cloudfront
Via: 1.1 ebc28fb0ad14691ee5d6c1a49f41b878.cloudfront.net (CloudFront)
X-Amz-Cf-Id: Ngsjm0gBg2y4cDFG4uwpAmaKaT6Dejh7oGlVDMewQrUaeW_SPst_Bw==

{
  "jobSummaryList": [{
    "jobId": "68f0c163-fbd4-44e6-9fd1-25b14a434786",
    "jobName": "example"
  }]
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

RegisterJobDefinition

Registers an AWS Batch job definition.

Request Syntax

```
POST /v1/registerjobdefinition HTTP/1.1  
Content-type: application/json
```

```
{  
  "containerProperties": {  
    "command": [ "string" ],  
    "environment": [  
      {  
        "name": "string",  
        "value": "string"  
      }  
    ],  
    "image": "string",  
    "jobRoleArn": "string",  
    "memory": number,  
    "mountPoints": [  
      {  
        "containerPath": "string",  
        "readOnly": boolean,  
        "sourceVolume": "string"  
      }  
    ],  
    "privileged": boolean,  
    "readonlyRootFilesystem": boolean,  
    "ulimits": [  
      {  
        "hardLimit": number,  
        "name": "string",  
        "softLimit": number  
      }  
    ],  
    "user": "string",  
    "vcpus": number,  
    "volumes": [  
      {  
        "host": {  
          "sourcePath": "string"  
        },  
        "name": "string"  
      }  
    ]  
  },  
  "jobDefinitionName": "string",  
  "parameters": {  
    "string" : "string"  
  },  
  "retryStrategy": {  
    "attempts": number  
  },  
  "timeout": {  
    "attemptDurationSeconds": number  
  },  
  "type": "string"  
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

containerProperties (p. 49)

An object with various properties specific for container-based jobs. This parameter is required if the `type` parameter is `container`.

Type: [ContainerProperties](#) (p. 88) object

Required: No

jobDefinitionName (p. 49)

The name of the job definition to register. Up to 128 letters (uppercase and lowercase), numbers, hyphens, and underscores are allowed.

Type: String

Required: Yes

parameters (p. 49)

Default parameter substitution placeholders to set in the job definition. Parameters are specified as a key-value pair mapping. Parameters in a `SubmitJob` request override any corresponding parameter defaults from the job definition.

Type: String to string map

Required: No

retryStrategy (p. 49)

The retry strategy to use for failed jobs that are submitted with this job definition. Any retry strategy that is specified during a [SubmitJob](#) (p. 55) operation overrides the retry strategy defined here. If a job is terminated due to a timeout, it is not retried.

Type: [RetryStrategy](#) (p. 106) object

Required: No

timeout (p. 49)

The timeout configuration for jobs that are submitted with this job definition, after which AWS Batch terminates your jobs if they have not finished. If a job is terminated due to a timeout, it is not retried. The minimum value for the timeout is 60 seconds. Any timeout configuration that is specified during a [SubmitJob](#) (p. 55) operation overrides the timeout configuration defined here. For more information, see [Job Timeouts](#) in the *Amazon Elastic Container Service Developer Guide*.

Type: [JobTimeout](#) (p. 103) object

Required: No

type (p. 49)

The type of job definition.

Type: String

Valid Values: `container`

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "jobDefinitionArn": "string",
  "jobDefinitionName": "string",
  "revision": number
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

jobDefinitionArn (p. 51)

The Amazon Resource Name (ARN) of the job definition.

Type: String

jobDefinitionName (p. 51)

The name of the job definition.

Type: String

revision (p. 51)

The revision of the job definition.

Type: Integer

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Examples

In the following example or examples, the Authorization header contents (`AUTHPARAMS`) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example registers a job definition for a simple container job.

Sample Request

```
POST /v1/registerjobdefinition HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 153
Authorization: AUTHPARAMS
X-Amz-Date: 20161128T215526Z
User-Agent: aws-cli/1.11.21 Python/2.7.12 Darwin/16.1.0 botocore/1.4.78

{
  "containerProperties": {
    "image": "busybox",
    "command": [
      "sleep",
      "10"
    ],
    "vcpus": 1,
    "memory": 128
  },
  "type": "container",
  "jobDefinitionName": "sleep10"
}
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 127
Connection: keep-alive
Date: Mon, 28 Nov 2016 21:55:27 GMT
x-amzn-RequestId: 5f0a08ec-b5b5-11e6-90ed-f72854f7eb90
X-Amzn-Trace-Id: Root=1-583ca7ce-bda1f338231616471b2efced
X-Cache: Miss from cloudfront
Via: 1.1 7a06af51e583997d8673ab89482dd45a.cloudfront.net (CloudFront)
X-Amz-Cf-Id: Y14HPNWWcKgm1U0wJPFfLBzLDvrMSdyuHo4GMi0oQwI0ukruLpi0nFw==

{
  "jobDefinitionName": "sleep10",
  "jobDefinitionArn": "arn:aws:batch:us-east-1:012345678910:job-definition/sleep10:1",
  "revision": 1
}
```

Example

This example registers a job definition for a simple container job with retries.

Sample Request

```
POST /v1/registerjobdefinition HTTP/1.1
```

```
Host: batch-beta.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 320
Authorization: AUTHPARAMS
X-Amz-Date: 20170327T145208Z
User-Agent: aws-cli/1.11.66 Python/2.7.10 Darwin/16.4.0 botocore/1.5.29

{
  "containerProperties": {
    "mountPoints": [],
    "image": "amazonlinux",
    "environment": [],
    "vcpus": 1,
    "command": [
      "/bin/bash",
      "-c",
      "exit $AWS_BATCH_JOB_ATTEMPT"
    ],
    "volumes": [],
    "memory": 2,
    "ulimits": []
  },
  "retryStrategy": {
    "attempts": 3
  },
  "jobDefinitionName": "EchoAttemptNumber",
  "parameters": {},
  "type": "container"
}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 27 Mar 2017 14:51:58 GMT
Content-Type: application/json
Content-Length: 147
Connection: keep-alive
x-amzn-RequestId: edbe4e89-12fc-11e7-a2c6-31d9fbe506fc
X-Amzn-Trace-Id: Root=1-58d9270e-56b3c6e2a91e74ef4c35a9c5
X-Cache: Miss from cloudfront
Via: 1.1 861b49a34b383ce3ac4ea8b7117b8953.cloudfront.net (CloudFront)
X-Amz-Cf-Id: l3zCxHtlIx4c1-RN2vkqIlpSbU9tsZNxfMSg6oIf700wg1BX0V5f_A==

{
  "jobDefinitionName": "EchoAttemptNumber",
  "jobDefinitionArn": "arn:aws:batch:us-east-1:012345678910:job-definition/
EchoAttemptNumber:1",
  "revision": 1
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)

- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

SubmitJob

Submits an AWS Batch job from a job definition. Parameters specified during [SubmitJob \(p. 55\)](#) override parameters defined in the job definition.

Request Syntax

```
POST /v1/submitjob HTTP/1.1
Content-type: application/json

{
  "arrayProperties": {
    "size": number
  },
  "containerOverrides": {
    "command": [ string ],
    "environment": [
      {
        "name": string,
        "value": string
      }
    ],
    "memory": number,
    "vcpus": number
  },
  "dependsOn": [
    {
      "jobId": string,
      "type": string
    }
  ],
  "jobDefinition": string,
  "jobName": string,
  "jobQueue": string,
  "parameters": {
    string : string
  },
  "retryStrategy": {
    "attempts": number
  },
  "timeout": {
    "attemptDurationSeconds": number
  }
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[arrayProperties \(p. 55\)](#)

The array properties for the submitted job, such as the size of the array. The array size can be between 2 and 10,000. If you specify array properties for a job, it becomes an array job. For more information, see [Array Jobs](#) in the *AWS Batch User Guide*.

Type: [ArrayProperties \(p. 71\)](#) object

Required: No

[containerOverrides \(p. 55\)](#)

A list of container overrides in JSON format that specify the name of a container in the specified job definition and the overrides it should receive. You can override the default command for a container (that is specified in the job definition or the Docker image) with a `command` override. You can also override existing environment variables (that are specified in the job definition or Docker image) on a container or add new environment variables to it with an `environment` override.

Type: [ContainerOverrides \(p. 87\)](#) object

Required: No

[dependsOn \(p. 55\)](#)

A list of dependencies for the job. A job can depend upon a maximum of 20 jobs. You can specify a `SEQUENTIAL` type dependency without specifying a job ID for array jobs so that each child array job completes sequentially, starting at index 0. You can also specify an `N_TO_N` type dependency with a job ID for array jobs so that each index child of this job must wait for the corresponding index child of each dependency to complete before it can begin.

Type: Array of [JobDependency \(p. 95\)](#) objects

Required: No

[jobDefinition \(p. 55\)](#)

The job definition used by this job. This value can be either a `name:revision` or the Amazon Resource Name (ARN) for the job definition.

Type: String

Required: Yes

[jobName \(p. 55\)](#)

The name of the job. The first character must be alphanumeric, and up to 128 letters (uppercase and lowercase), numbers, hyphens, and underscores are allowed.

Type: String

Required: Yes

[jobQueue \(p. 55\)](#)

The job queue into which the job is submitted. You can specify either the name or the Amazon Resource Name (ARN) of the queue.

Type: String

Required: Yes

[parameters \(p. 55\)](#)

Additional parameters passed to the job that replace parameter substitution placeholders that are set in the job definition. Parameters are specified as a key and value pair mapping. Parameters in a `SubmitJob` request override any corresponding parameter defaults from the job definition.

Type: String to string map

Required: No

[retryStrategy \(p. 55\)](#)

The retry strategy to use for failed jobs from this [SubmitJob \(p. 55\)](#) operation. When a retry strategy is specified here, it overrides the retry strategy defined in the job definition.

Type: [RetryStrategy \(p. 106\)](#) object

Required: No

[timeout \(p. 55\)](#)

The timeout configuration for this [SubmitJob \(p. 55\)](#) operation. You can specify a timeout duration after which AWS Batch terminates your jobs if they have not finished. If a job is terminated due to a timeout, it is not retried. The minimum value for the timeout is 60 seconds. This configuration overrides any timeout configuration specified in the job definition. For array jobs, child jobs have the same timeout configuration as the parent job. For more information, see [Job Timeouts](#) in the *Amazon Elastic Container Service Developer Guide*.

Type: [JobTimeout \(p. 103\)](#) object

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "jobId": "string",
  "jobName": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[jobId \(p. 57\)](#)

The unique identifier for the job.

Type: String

[jobName \(p. 57\)](#)

The name of the job.

Type: String

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Example

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example submits a simple container job called `example` to the `HighPriority` job queue.

Sample Request

```
POST /v1/submitjob HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 78
Authorization: AUTHPARAMS
X-Amz-Date: 20161129T201034Z
User-Agent: aws-cli/1.11.22 Python/2.7.12 Darwin/16.1.0 botocore/1.4.79

{
  "jobName": "example",
  "jobQueue": "HighPriority",
  "jobDefinition": "sleep60"
}
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 82
Connection: keep-alive
Date: Tue, 29 Nov 2016 20:10:34 GMT
x-amzn-RequestId: e2e433cf-b66f-11e6-8321-7fedcfd554e5
X-Amzn-Trace-Id: Root=1-583de0ba-d26be22c375ed3416b2e18b7
X-Cache: Miss from cloudfront
Via: 1.1 6ba12aef47e3e7677e084594bfce5e1.cloudfront.net (CloudFront)
X-Amz-Cf-Id: QoWCvISTvYBbaP2hAoxC8_TWTl2JN-kNASYCjCJ5HztN0e1OuzvpSA==

{
  "jobName": "example",
  "jobId": "876da822-4198-45f2-a252-6cea32512ea8"
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

TerminateJob

Terminates a job in a job queue. Jobs that are in the `STARTING` or `RUNNING` state are terminated, which causes them to transition to `FAILED`. Jobs that have not progressed to the `STARTING` state are cancelled.

Request Syntax

```
POST /v1/terminatejob HTTP/1.1
Content-type: application/json

{
  "jobId": "string",
  "reason": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[jobId \(p. 60\)](#)

The AWS Batch job ID of the job to terminate.

Type: String

Required: Yes

[reason \(p. 60\)](#)

A message to attach to the job that explains the reason for canceling it. This message is returned by future [DescribeJobs \(p. 39\)](#) operations on the job. This message is also recorded in the AWS Batch activity logs.

Type: String

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Example

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example terminates a job with the specified job ID.

Sample Request

```
POST /v1/terminatejob HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 79
Authorization: AUTHPARAMS
X-Amz-Date: 20161129T202905Z
User-Agent: aws-cli/1.11.22 Python/2.7.12 Darwin/16.1.0 botocore/1.4.79

{
  "reason": "Terminating job.",
  "jobId": "61e743ed-35e4-48da-b2de-5c8333821c84"
}
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 2
Connection: keep-alive
Date: Tue, 29 Nov 2016 20:29:06 GMT
x-amzn-RequestId: 795eee02-b672-11e6-8460-6d7ce4bf85d3
X-Amzn-Trace-Id: Root=1-583de512-8c218eec31cef60008ef7d93
X-Cache: Miss from cloudfront
Via: 1.1 16d2657cebef5191828b055567b4efeb.cloudfront.net (CloudFront)
X-Amz-Cf-Id: 681NTs_bPulMwja2HekWMwngcUzx2a8w_oaG27W0L4Pjct7W1T-Fvw==
```

```
{ }
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

UpdateComputeEnvironment

Updates an AWS Batch compute environment.

Request Syntax

```
POST /v1/updatecomputeenvironment HTTP/1.1
Content-type: application/json

{
  "computeEnvironment": "string",
  "computeResources": {
    "desiredvCpus": number,
    "maxvCpus": number,
    "minvCpus": number
  },
  "serviceRole": "string",
  "state": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

computeEnvironment (p. 63)

The name or full Amazon Resource Name (ARN) of the compute environment to update.

Type: String

Required: Yes

computeResources (p. 63)

Details of the compute resources managed by the compute environment. Required for a managed compute environment.

Type: [ComputeResourceUpdate](#) (p. 83) object

Required: No

serviceRole (p. 63)

The full Amazon Resource Name (ARN) of the IAM role that allows AWS Batch to make calls to other AWS services on your behalf.

If your specified role has a path other than /, then you must either specify the full role ARN (this is recommended) or prefix the role name with the path.

Note

Depending on how you created your AWS Batch service role, its ARN may contain the `service-role` path prefix. When you only specify the name of the service role, AWS Batch assumes that your ARN does not use the `service-role` path prefix. Because of this, we recommend that you specify the full ARN of your service role when you create compute environments.

Type: String

Required: No

[state \(p. 63\)](#)

The state of the compute environment. Compute environments in the `ENABLED` state can accept jobs from a queue and scale in or out automatically based on the workload demand of its associated queues.

Type: String

Valid Values: `ENABLED` | `DISABLED`

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "computeEnvironmentArn": "string",
  "computeEnvironmentName": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[computeEnvironmentArn \(p. 64\)](#)

The Amazon Resource Name (ARN) of the compute environment.

Type: String

[computeEnvironmentName \(p. 64\)](#)

The name of compute environment.

Type: String

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Example

In the following example or examples, the Authorization header contents (AUTHPARAMS) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example disables the P3OnDemand compute environment so it can be deleted.

Sample Request

```
POST /v1/updatecomputeenvironment HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 57
Authorization: AUTHPARAMS
X-Amz-Date: 20161128T194248Z
User-Agent: aws-cli/1.11.21 Python/2.7.12 Darwin/16.1.0 botocore/1.4.78

{
  "computeEnvironment": "P3OnDemand",
  "state": "DISABLED"
}
```

Sample Response

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 133
Connection: keep-alive
Date: Mon, 28 Nov 2016 19:42:49 GMT
x-amzn-RequestId: d7d41aba-b5a2-11e6-bbde-956d603f3192
X-Amzn-Trace-Id: Root=1-583c88b9-c30dd12f24398eef8bd95ed7
X-Cache: Miss from cloudfront
Via: 1.1 7f3f42df8af148df1f9f1ee7175987ad.cloudfront.net (CloudFront)
X-Amz-Cf-Id: uxJn0P7cg_1RTxOs15FkCItWfmCeniKMZdXlFWaOfPfqATHw3j-AA==

{
  "computeEnvironmentName": "P3OnDemand",
  "computeEnvironmentArn": "arn:aws:batch:us-east-1:012345678910:compute-environment/
P3OnDemand"
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)

- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

UpdateJobQueue

Updates a job queue.

Request Syntax

```
POST /v1/updatejobqueue HTTP/1.1
Content-type: application/json

{
  "computeEnvironmentOrder": [
    {
      "computeEnvironment": "string",
      "order": number
    }
  ],
  "jobQueue": "string",
  "priority": number,
  "state": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

computeEnvironmentOrder (p. 67)

Details the set of compute environments mapped to a job queue and their order relative to each other. This is one of the parameters used by the job scheduler to determine which compute environment should execute a given job.

Type: Array of [ComputeEnvironmentOrder](#) (p. 79) objects

Required: No

jobQueue (p. 67)

The name or the Amazon Resource Name (ARN) of the job queue.

Type: String

Required: Yes

priority (p. 67)

The priority of the job queue. Job queues with a higher priority (or a higher integer value for the `priority` parameter) are evaluated first when associated with same compute environment. Priority is determined in descending order, for example, a job queue with a priority value of 10 is given scheduling preference over a job queue with a priority value of 1.

Type: Integer

Required: No

state (p. 67)

Describes the queue's ability to accept new jobs.

Type: String

Valid Values: ENABLED | DISABLED

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "jobQueueArn": "string",
  "jobQueueName": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[jobQueueArn \(p. 68\)](#)

The Amazon Resource Name (ARN) of the job queue.

Type: String

[jobQueueName \(p. 68\)](#)

The name of the job queue.

Type: String

Errors

ClientException

These errors are usually caused by a client action, such as using an action or resource on behalf of a user that doesn't have permissions to use the action or resource, or specifying an identifier that is not valid.

HTTP Status Code: 400

ServerException

These errors are usually caused by a server issue.

HTTP Status Code: 500

Example

In the following example or examples, the Authorization header contents (`AUTHPARAMS`) must be replaced with an AWS Signature Version 4 signature. For more information about creating these signatures, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

You only need to learn how to sign HTTP requests if you intend to manually create them. When you use the [AWS Command Line Interface \(AWS CLI\)](#) or one of the [AWS SDKs](#) to make requests to AWS, these tools automatically sign the requests for you with the access key that you specify when you configure the tools. When you use these tools, you don't need to learn how to sign requests yourself.

Example

This example disables a job queue so that it can be deleted.

Sample Request

```
POST /v1/updatejobqueue HTTP/1.1
Host: batch.us-east-1.amazonaws.com
Accept-Encoding: identity
Content-Length: 42
Authorization: AUTHPARAMS
X-Amz-Date: 20161128T201802Z
User-Agent: aws-cli/1.11.21 Python/2.7.12 Darwin/16.1.0 botocore/1.4.78

{
  "state": "DISABLED",
  "jobQueue": "GPGPU"
}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 28 Nov 2016 20:18:03 GMT
Content-Type: application/json
Content-Length: 93
Connection: keep-alive
x-amzn-RequestId: c3ad778b-b5a7-11e6-8ed7-ed4d5ff654a
X-Amzn-Trace-Id: Root=1-583c90fa-b6f9f9acb4b6f8bc81725c75
X-Cache: Miss from cloudfront
Via: 1.1 17de248e6d780f737234d37cc490dbe3.cloudfront.net (CloudFront)
X-Amz-Cf-Id: aVju0hE8eLpjSF18Y3fOuxgOZXdigQLcDMw00plxnynw0dEsOsEgw==

{
  "jobQueueName": "GPGPU",
  "jobQueueArn": "arn:aws:batch:us-east-1:012345678910:job-queue/GPGPU"
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

Data Types

The AWS Batch API contains several data types that various actions use. This section describes each data type in detail.

Note

The order of each element in a data type structure is not guaranteed. Applications should not assume a particular order.

The following data types are supported:

- [ArrayProperties](#) (p. 71)
- [ArrayPropertiesDetail](#) (p. 72)
- [ArrayPropertiesSummary](#) (p. 73)
- [AttemptContainerDetail](#) (p. 74)
- [AttemptDetail](#) (p. 76)
- [ComputeEnvironmentDetail](#) (p. 77)
- [ComputeEnvironmentOrder](#) (p. 79)
- [ComputeResource](#) (p. 80)
- [ComputeResourceUpdate](#) (p. 83)
- [ContainerDetail](#) (p. 84)
- [ContainerOverrides](#) (p. 87)
- [ContainerProperties](#) (p. 88)
- [ContainerSummary](#) (p. 91)
- [Host](#) (p. 92)
- [JobDefinition](#) (p. 93)
- [JobDependency](#) (p. 95)
- [JobDetail](#) (p. 96)
- [JobQueueDetail](#) (p. 99)
- [JobSummary](#) (p. 101)
- [JobTimeout](#) (p. 103)
- [KeyValuePair](#) (p. 104)
- [MountPoint](#) (p. 105)
- [RetryStrategy](#) (p. 106)
- [Ulimit](#) (p. 107)
- [Volume](#) (p. 108)

ArrayProperties

An object representing an AWS Batch array job.

Contents

size

The size of the array job.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ArrayPropertiesDetail

An object representing the array properties of a job.

Contents

index

The job index within the array that is associated with this job. This parameter is returned for array job children.

Type: Integer

Required: No

size

The size of the array job. This parameter is returned for parent array jobs.

Type: Integer

Required: No

statusSummary

A summary of the number of array job children in each available job status. This parameter is returned for parent array jobs.

Type: String to integer map

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ArrayPropertiesSummary

An object representing the array properties of a job.

Contents

index

The job index within the array that is associated with this job. This parameter is returned for children of array jobs.

Type: Integer

Required: No

size

The size of the array job. This parameter is returned for parent array jobs.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

AttemptContainerDetail

An object representing the details of a container that is part of a job attempt.

Contents

containerInstanceArn

The Amazon Resource Name (ARN) of the Amazon ECS container instance that hosts the job attempt.

Type: String

Required: No

exitCode

The exit code for the job attempt. A non-zero exit code is considered a failure.

Type: Integer

Required: No

logStreamName

The name of the CloudWatch Logs log stream associated with the container. The log group for AWS Batch jobs is `/aws/batch/job`. Each container attempt receives a log stream name when they reach the `RUNNING` status.

Type: String

Required: No

reason

A short (255 max characters) human-readable string to provide additional details about a running or stopped container.

Type: String

Required: No

taskArn

The Amazon Resource Name (ARN) of the Amazon ECS task that is associated with the job attempt. Each container attempt receives a task ARN when they reach the `STARTING` status.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

AttemptDetail

An object representing a job attempt.

Contents

container

Details about the container in this job attempt.

Type: [AttemptContainerDetail](#) (p. 74) object

Required: No

startedAt

The Unix time stamp (in seconds and milliseconds) for when the attempt was started (when the attempt transitioned from the `STARTING` state to the `RUNNING` state).

Type: Long

Required: No

statusReason

A short, human-readable string to provide additional details about the current status of the job attempt.

Type: String

Required: No

stoppedAt

The Unix time stamp (in seconds and milliseconds) for when the attempt was stopped (when the attempt transitioned from the `RUNNING` state to a terminal state, such as `SUCCEEDED` or `FAILED`).

Type: Long

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ComputeEnvironmentDetail

An object representing an AWS Batch compute environment.

Contents

computeEnvironmentArn

The Amazon Resource Name (ARN) of the compute environment.

Type: String

Required: Yes

computeEnvironmentName

The name of the compute environment.

Type: String

Required: Yes

computeResources

The compute resources defined for the compute environment.

Type: [ComputeResource \(p. 80\)](#) object

Required: No

ecsClusterArn

The Amazon Resource Name (ARN) of the underlying Amazon ECS cluster used by the compute environment.

Type: String

Required: Yes

serviceRole

The service role associated with the compute environment that allows AWS Batch to make calls to AWS API operations on your behalf.

Type: String

Required: No

state

The state of the compute environment. The valid values are `ENABLED` or `DISABLED`. An `ENABLED` state indicates that you can register instances with the compute environment and that the associated instances can accept jobs.

Type: String

Valid Values: `ENABLED` | `DISABLED`

Required: No

status

The current status of the compute environment (for example, `CREATING` or `VALID`).

Type: String

Valid Values: CREATING | UPDATING | DELETING | DELETED | VALID | INVALID

Required: No

statusReason

A short, human-readable string to provide additional details about the current status of the compute environment.

Type: String

Required: No

type

The type of the compute environment.

Type: String

Valid Values: MANAGED | UNMANAGED

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ComputeEnvironmentOrder

The order in which compute environments are tried for job placement within a queue. Compute environments are tried in ascending order. For example, if two compute environments are associated with a job queue, the compute environment with a lower order integer value is tried for job placement first.

Contents

computeEnvironment

The Amazon Resource Name (ARN) of the compute environment.

Type: String

Required: Yes

order

The order of the compute environment.

Type: Integer

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ComputeResource

An object representing an AWS Batch compute resource.

Contents

bidPercentage

The minimum percentage that a Spot Instance price must be when compared with the On-Demand price for that instance type before instances are launched. For example, if your `bidPercentage` is 20, then the Spot price must be below 20% of the current On-Demand price for that EC2 instance.

Type: Integer

Required: No

desiredvCpus

The desired number of EC2 vCPUS in the compute environment.

Type: Integer

Required: No

ec2KeyPair

The EC2 key pair that is used for instances launched in the compute environment.

Type: String

Required: No

imageId

The Amazon Machine Image (AMI) ID used for instances launched in the compute environment.

Type: String

Required: No

instanceRole

The Amazon ECS instance profile applied to Amazon EC2 instances in a compute environment. You can specify the short name or full Amazon Resource Name (ARN) of an instance profile. For example, `ecsInstanceRole` or `arn:aws:iam::<aws_account_id>:instance-profile/ecsInstanceRole`. For more information, see [Amazon ECS Instance Role](#) in the *AWS Batch User Guide*.

Type: String

Required: Yes

instanceTypes

The instances types that may be launched. You can specify instance families to launch any instance type within those families (for example, `c4` or `p3`), or you can specify specific sizes within a family (such as `c4.8xlarge`). You can also choose `optimal` to pick instance types (from the latest C, M, and R instance families) on the fly that match the demand of your job queues.

Type: Array of strings

Required: Yes

maxvCpus

The maximum number of EC2 vCPUs that an environment can reach.

Type: Integer

Required: Yes

minvCpus

The minimum number of EC2 vCPUs that an environment should maintain.

Type: Integer

Required: Yes

securityGroupIds

The EC2 security group that is associated with instances launched in the compute environment.

Type: Array of strings

Required: Yes

spotIamFleetRole

The Amazon Resource Name (ARN) of the Amazon EC2 Spot Fleet IAM role applied to a `SPOT` compute environment.

Type: String

Required: No

subnets

The VPC subnets into which the compute resources are launched.

Type: Array of strings

Required: Yes

tags

Key-value pair tags to be applied to resources that are launched in the compute environment.

Type: String to string map

Required: No

type

The type of compute environment.

Type: String

Valid Values: `EC2` | `SPOT`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ComputeResourceUpdate

An object representing the attributes of a compute environment that can be updated.

Contents

desiredvCpus

The desired number of EC2 vCPUS in the compute environment.

Type: Integer

Required: No

maxvCpus

The maximum number of EC2 vCPUs that an environment can reach.

Type: Integer

Required: No

minvCpus

The minimum number of EC2 vCPUs that an environment should maintain.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ContainerDetail

An object representing the details of a container that is part of a job.

Contents

command

The command that is passed to the container.

Type: Array of strings

Required: No

containerInstanceArn

The Amazon Resource Name (ARN) of the container instance on which the container is running.

Type: String

Required: No

environment

The environment variables to pass to a container.

Note

Environment variables must not start with `AWS_BATCH`; this naming convention is reserved for variables that are set by the AWS Batch service.

Type: Array of [KeyValuePair](#) (p. 104) objects

Required: No

exitCode

The exit code to return upon completion.

Type: Integer

Required: No

image

The image used to start the container.

Type: String

Required: No

jobRoleArn

The Amazon Resource Name (ARN) associated with the job upon execution.

Type: String

Required: No

logStreamName

The name of the CloudWatch Logs log stream associated with the container. The log group for AWS Batch jobs is `/aws/batch/job`. Each container attempt receives a log stream name when they reach the `RUNNING` status.

Type: String

Required: No

memory

The number of MiB of memory reserved for the job.

Type: Integer

Required: No

mountPoints

The mount points for data volumes in your container.

Type: Array of [MountPoint \(p. 105\)](#) objects

Required: No

privileged

When this parameter is true, the container is given elevated privileges on the host container instance (similar to the `root` user).

Type: Boolean

Required: No

readonlyRootFilesystem

When this parameter is true, the container is given read-only access to its root file system.

Type: Boolean

Required: No

reason

A short (255 max characters) human-readable string to provide additional details about a running or stopped container.

Type: String

Required: No

taskArn

The Amazon Resource Name (ARN) of the Amazon ECS task that is associated with the container job. Each container attempt receives a task ARN when they reach the `STARTING` status.

Type: String

Required: No

ulimits

A list of `ulimit` values to set in the container.

Type: Array of [Ulimit \(p. 107\)](#) objects

Required: No

user

The user name to use inside the container.

Type: String

Required: No

vcpus

The number of VCPUs allocated for the job.

Type: Integer

Required: No

volumes

A list of volumes associated with the job.

Type: Array of [Volume \(p. 108\)](#) objects

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ContainerOverrides

The overrides that should be sent to a container.

Contents

command

The command to send to the container that overrides the default command from the Docker image or the job definition.

Type: Array of strings

Required: No

environment

The environment variables to send to the container. You can add new environment variables, which are added to the container at launch, or you can override the existing environment variables from the Docker image or the job definition.

Note

Environment variables must not start with `AWS_BATCH`; this naming convention is reserved for variables that are set by the AWS Batch service.

Type: Array of [KeyValuePair \(p. 104\)](#) objects

Required: No

memory

The number of MiB of memory reserved for the job. This value overrides the value set in the job definition.

Type: Integer

Required: No

vcpus

The number of vCPUs to reserve for the container. This value overrides the value set in the job definition.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ContainerProperties

Container properties are used in job definitions to describe the container that is launched as part of a job.

Contents

command

The command that is passed to the container. This parameter maps to `Cmd` in the [Create a container](#) section of the [Docker Remote API](#) and the `COMMAND` parameter to `docker run`. For more information, see <https://docs.docker.com/engine/reference/builder/#cmd>.

Type: Array of strings

Required: No

environment

The environment variables to pass to a container. This parameter maps to `Env` in the [Create a container](#) section of the [Docker Remote API](#) and the `--env` option to `docker run`.

Important

We do not recommend using plaintext environment variables for sensitive information, such as credential data.

Note

Environment variables must not start with `AWS_BATCH`; this naming convention is reserved for variables that are set by the AWS Batch service.

Type: Array of [KeyValuePair \(p. 104\)](#) objects

Required: No

image

The image used to start a container. This string is passed directly to the Docker daemon. Images in the Docker Hub registry are available by default. Other repositories are specified with `repository-url/image:tag`. Up to 255 letters (uppercase and lowercase), numbers, hyphens, underscores, colons, periods, forward slashes, and number signs are allowed. This parameter maps to `Image` in the [Create a container](#) section of the [Docker Remote API](#) and the `IMAGE` parameter of `docker run`.

- Images in Amazon ECR repositories use the full registry and repository URI (for example, `012345678910.dkr.ecr.<region-name>.amazonaws.com/<repository-name>`).
- Images in official repositories on Docker Hub use a single name (for example, `ubuntu` or `mongo`).
- Images in other repositories on Docker Hub are qualified with an organization name (for example, `amazon/amazon-ecs-agent`).
- Images in other online repositories are qualified further by a domain name (for example, `quay.io/assemblyline/ubuntu`).

Type: String

Required: Yes

jobRoleArn

The Amazon Resource Name (ARN) of the IAM role that the container can assume for AWS permissions.

Type: String

Required: No

memory

The hard limit (in MiB) of memory to present to the container. If your container attempts to exceed the memory specified here, the container is killed. This parameter maps to `Memory` in the [Create a container](#) section of the [Docker Remote API](#) and the `--memory` option to `docker run`. You must specify at least 4 MiB of memory for a job.

Note

If you are trying to maximize your resource utilization by providing your jobs as much memory as possible for a particular instance type, see [Memory Management](#) in the *AWS Batch User Guide*.

Type: Integer

Required: Yes

mountPoints

The mount points for data volumes in your container. This parameter maps to `Volumes` in the [Create a container](#) section of the [Docker Remote API](#) and the `--volume` option to `docker run`.

Type: Array of [MountPoint \(p. 105\)](#) objects

Required: No

privileged

When this parameter is true, the container is given elevated privileges on the host container instance (similar to the `root` user). This parameter maps to `Privileged` in the [Create a container](#) section of the [Docker Remote API](#) and the `--privileged` option to `docker run`.

Type: Boolean

Required: No

readonlyRootFilesystem

When this parameter is true, the container is given read-only access to its root file system. This parameter maps to `ReadOnlyRootfs` in the [Create a container](#) section of the [Docker Remote API](#) and the `--read-only` option to `docker run`.

Type: Boolean

Required: No

ulimits

A list of `ulimits` to set in the container. This parameter maps to `Ulimits` in the [Create a container](#) section of the [Docker Remote API](#) and the `--ulimit` option to `docker run`.

Type: Array of [Ulimit \(p. 107\)](#) objects

Required: No

user

The user name to use inside the container. This parameter maps to `User` in the [Create a container](#) section of the [Docker Remote API](#) and the `--user` option to `docker run`.

Type: String

Required: No

vcpus

The number of vCPUs reserved for the container. This parameter maps to `CpuShares` in the [Create a container](#) section of the [Docker Remote API](#) and the `--cpu-shares` option to `docker run`. Each vCPU is equivalent to 1,024 CPU shares. You must specify at least one vCPU.

Type: Integer

Required: Yes

volumes

A list of data volumes used in a job.

Type: Array of [Volume \(p. 108\)](#) objects

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ContainerSummary

An object representing summary details of a container within a job.

Contents

exitCode

The exit code to return upon completion.

Type: Integer

Required: No

reason

A short (255 max characters) human-readable string to provide additional details about a running or stopped container.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

Host

The contents of the `host` parameter determine whether your data volume persists on the host container instance and where it is stored. If the `host` parameter is empty, then the Docker daemon assigns a host path for your data volume, but the data is not guaranteed to persist after the containers associated with it stop running.

Contents

sourcePath

The path on the host container instance that is presented to the container. If this parameter is empty, then the Docker daemon has assigned a host path for you. If the `host` parameter contains a `sourcePath` file location, then the data volume persists at the specified location on the host container instance until you delete it manually. If the `sourcePath` value does not exist on the host container instance, the Docker daemon creates it. If the location does exist, the contents of the source path folder are exported.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

JobDefinition

An object representing an AWS Batch job definition.

Contents

containerProperties

An object with various properties specific to container-based jobs.

Type: [ContainerProperties \(p. 88\)](#) object

Required: No

jobDefinitionArn

The Amazon Resource Name (ARN) for the job definition.

Type: String

Required: Yes

jobDefinitionName

The name of the job definition.

Type: String

Required: Yes

parameters

Default parameters or parameter substitution placeholders that are set in the job definition. Parameters are specified as a key-value pair mapping. Parameters in a `SubmitJob` request override any corresponding parameter defaults from the job definition.

Type: String to string map

Required: No

retryStrategy

The retry strategy to use for failed jobs that are submitted with this job definition.

Type: [RetryStrategy \(p. 106\)](#) object

Required: No

revision

The revision of the job definition.

Type: Integer

Required: Yes

status

The status of the job definition.

Type: String

Required: No

timeout

The timeout configuration for jobs that are submitted with this job definition. You can specify a timeout duration after which AWS Batch terminates your jobs if they have not finished.

Type: [JobTimeout \(p. 103\)](#) object

Required: No

type

The type of job definition.

Type: String

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

JobDependency

An object representing an AWS Batch job dependency.

Contents

jobId

The job ID of the AWS Batch job associated with this dependency.

Type: String

Required: No

type

The type of the job dependency.

Type: String

Valid Values: `N_TO_N` | `SEQUENTIAL`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

JobDetail

An object representing an AWS Batch job.

Contents

arrayProperties

The array properties of the job, if it is an array job.

Type: [ArrayPropertiesDetail \(p. 72\)](#) object

Required: No

attempts

A list of job attempts associated with this job.

Type: Array of [AttemptDetail \(p. 76\)](#) objects

Required: No

container

An object representing the details of the container that is associated with the job.

Type: [ContainerDetail \(p. 84\)](#) object

Required: No

createdAt

The Unix time stamp (in seconds and milliseconds) for when the job was created. For non-array jobs and parent array jobs, this is when the job entered the `SUBMITTED` state (at the time [SubmitJob \(p. 55\)](#) was called). For array child jobs, this is when the child job was spawned by its parent and entered the `PENDING` state.

Type: Long

Required: No

dependsOn

A list of job names or IDs on which this job depends.

Type: Array of [JobDependency \(p. 95\)](#) objects

Required: No

jobDefinition

The job definition that is used by this job.

Type: String

Required: Yes

jobId

The ID for the job.

Type: String

Required: Yes

jobName

The name of the job.

Type: String

Required: Yes

jobQueue

The Amazon Resource Name (ARN) of the job queue with which the job is associated.

Type: String

Required: Yes

parameters

Additional parameters passed to the job that replace parameter substitution placeholders or override any corresponding parameter defaults from the job definition.

Type: String to string map

Required: No

retryStrategy

The retry strategy to use for this job if an attempt fails.

Type: [RetryStrategy \(p. 106\)](#) object

Required: No

startedAt

The Unix time stamp (in seconds and milliseconds) for when the job was started (when the job transitioned from the `STARTING` state to the `RUNNING` state).

Type: Long

Required: Yes

status

The current status for the job.

Type: String

Valid Values: `SUBMITTED` | `PENDING` | `RUNNABLE` | `STARTING` | `RUNNING` | `SUCCEEDED` | `FAILED`

Required: Yes

statusReason

A short, human-readable string to provide additional details about the current status of the job.

Type: String

Required: No

stoppedAt

The Unix time stamp (in seconds and milliseconds) for when the job was stopped (when the job transitioned from the `RUNNING` state to a terminal state, such as `SUCCEEDED` or `FAILED`).

Type: Long

Required: No

timeout

The timeout configuration for the job.

Type: [JobTimeout \(p. 103\)](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

JobQueueDetail

An object representing the details of an AWS Batch job queue.

Contents

computeEnvironmentOrder

The compute environments that are attached to the job queue and the order in which job placement is preferred. Compute environments are selected for job placement in ascending order.

Type: Array of [ComputeEnvironmentOrder](#) (p. 79) objects

Required: Yes

jobQueueArn

The Amazon Resource Name (ARN) of the job queue.

Type: String

Required: Yes

jobQueueName

The name of the job queue.

Type: String

Required: Yes

priority

The priority of the job queue.

Type: Integer

Required: Yes

state

Describes the ability of the queue to accept new jobs.

Type: String

Valid Values: `ENABLED` | `DISABLED`

Required: Yes

status

The status of the job queue (for example, `CREATING` or `VALID`).

Type: String

Valid Values: `CREATING` | `UPDATING` | `DELETING` | `DELETED` | `VALID` | `INVALID`

Required: No

statusReason

A short, human-readable string to provide additional details about the current status of the job queue.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

JobSummary

An object representing summary details of a job.

Contents

arrayProperties

The array properties of the job, if it is an array job.

Type: [ArrayPropertiesSummary \(p. 73\)](#) object

Required: No

container

An object representing the details of the container that is associated with the job.

Type: [ContainerSummary \(p. 91\)](#) object

Required: No

createdAt

The Unix time stamp for when the job was created. For non-array jobs and parent array jobs, this is when the job entered the `SUBMITTED` state (at the time [SubmitJob \(p. 55\)](#) was called). For array child jobs, this is when the child job was spawned by its parent and entered the `PENDING` state.

Type: Long

Required: No

jobId

The ID of the job.

Type: String

Required: Yes

jobName

The name of the job.

Type: String

Required: Yes

startedAt

The Unix time stamp for when the job was started (when the job transitioned from the `STARTING` state to the `RUNNING` state).

Type: Long

Required: No

status

The current status for the job.

Type: String

Valid Values: SUBMITTED | PENDING | RUNNABLE | STARTING | RUNNING | SUCCEEDED | FAILED

Required: No

statusReason

A short, human-readable string to provide additional details about the current status of the job.

Type: String

Required: No

stoppedAt

The Unix time stamp for when the job was stopped (when the job transitioned from the RUNNING state to a terminal state, such as SUCCEEDED or FAILED).

Type: Long

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

JobTimeout

An object representing a job timeout configuration.

Contents

attemptDurationSeconds

The time duration in seconds (measured from the job attempt's `startedAt` timestamp) after which AWS Batch terminates your jobs if they have not finished.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

KeyValuePair

A key-value pair object.

Contents

name

The name of the key-value pair. For environment variables, this is the name of the environment variable.

Type: String

Required: No

value

The value of the key-value pair. For environment variables, this is the value of the environment variable.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

MountPoint

Details on a Docker volume mount point that is used in a job's container properties.

Contents

containerPath

The path on the container at which to mount the host volume.

Type: String

Required: No

readOnly

If this value is `true`, the container has read-only access to the volume; otherwise, the container can write to the volume. The default value is `false`.

Type: Boolean

Required: No

sourceVolume

The name of the volume to mount.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

RetryStrategy

The retry strategy associated with a job.

Contents

attempts

The number of times to move a job to the `RUNNABLE` status. You may specify between 1 and 10 attempts. If the value of `attempts` is greater than one, the job is retried if it fails until it has moved to `RUNNABLE` that many times.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

Ulimit

The `ulimit` settings to pass to the container.

Contents

hardLimit

The hard limit for the `ulimit` type.

Type: Integer

Required: Yes

name

The type of the `ulimit`.

Type: String

Required: Yes

softLimit

The soft limit for the `ulimit` type.

Type: Integer

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

Volume

A data volume used in a job's container properties.

Contents

host

The contents of the `host` parameter determine whether your data volume persists on the host container instance and where it is stored. If the `host` parameter is empty, then the Docker daemon assigns a host path for your data volume. However, the data is not guaranteed to persist after the containers associated with it stop running.

Type: [Host \(p. 92\)](#) object

Required: No

name

The name of the volume. Up to 255 letters (uppercase and lowercase), numbers, hyphens, and underscores are allowed. This name is referenced in the `sourceVolume` parameter of container definition `mountPoints`.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

Common Parameters

The following list contains the parameters that all actions use for signing Signature Version 4 requests with a query string. Any action-specific parameters are listed in the topic for that action. For more information about Signature Version 4, see [Signature Version 4 Signing Process](#) in the *Amazon Web Services General Reference*.

Action

The action to be performed.

Type: string

Required: Yes

Version

The API version that the request is written for, expressed in the format YYYY-MM-DD.

Type: string

Required: Yes

X-Amz-Algorithm

The hash algorithm that you used to create the request signature.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Valid Values: `AWS4-HMAC-SHA256`

Required: Conditional

X-Amz-Credential

The credential scope value, which is a string that includes your access key, the date, the region you are targeting, the service you are requesting, and a termination string ("aws4_request"). The value is expressed in the following format: `access_key/YYYYMMDD/region/service/aws4_request`.

For more information, see [Task 2: Create a String to Sign for Signature Version 4](#) in the *Amazon Web Services General Reference*.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

X-Amz-Date

The date that is used to create the signature. The format must be ISO 8601 basic format (YYYYMMDD'THHMMSS'Z'). For example, the following date time is a valid X-Amz-Date value: `20120325T120000Z`.

Condition: X-Amz-Date is optional for all requests; it can be used to override the date used for signing requests. If the Date header is specified in the ISO 8601 basic format, X-Amz-Date is

not required. When X-Amz-Date is used, it always overrides the value of the Date header. For more information, see [Handling Dates in Signature Version 4](#) in the *Amazon Web Services General Reference*.

Type: string

Required: Conditional

X-Amz-Security-Token

The temporary security token that was obtained through a call to AWS Security Token Service (AWS STS). For a list of services that support temporary security credentials from AWS Security Token Service, go to [AWS Services That Work with IAM](#) in the *IAM User Guide*.

Condition: If you're using temporary security credentials from the AWS Security Token Service, you must include the security token.

Type: string

Required: Conditional

X-Amz-Signature

Specifies the hex-encoded signature that was calculated from the string to sign and the derived signing key.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

X-Amz-SignedHeaders

Specifies all the HTTP headers that were included as part of the canonical request. For more information about specifying signed headers, see [Task 1: Create a Canonical Request For Signature Version 4](#) in the *Amazon Web Services General Reference*.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional