



API Reference

AWS Compute Optimizer



API Version 2019-11-01

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

AWS Compute Optimizer: API Reference

Copyright © 2025 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
DeleteRecommendationPreferences	4
Request Syntax	4
Request Parameters	4
Response Elements	5
Errors	5
See Also	6
DescribeRecommendationExportJobs	8
Request Syntax	8
Request Parameters	8
Response Syntax	9
Response Elements	10
Errors	10
See Also	11
ExportAutoScalingGroupRecommendations	13
Request Syntax	13
Request Parameters	13
Response Syntax	17
Response Elements	17
Errors	18
See Also	19
ExportEBSVolumeRecommendations	20
Request Syntax	20
Request Parameters	20
Response Syntax	23
Response Elements	23
Errors	24
See Also	25
ExportEC2InstanceRecommendations	26
Request Syntax	26
Request Parameters	26
Response Syntax	30
Response Elements	30

Errors	31
See Also	32
ExportECSServiceRecommendations	33
Request Syntax	33
Request Parameters	33
Response Syntax	36
Response Elements	36
Errors	37
See Also	38
ExportIdleRecommendations	39
Request Syntax	39
Request Parameters	39
Response Syntax	42
Response Elements	42
Errors	42
See Also	44
ExportLambdaFunctionRecommendations	45
Request Syntax	45
Request Parameters	45
Response Syntax	48
Response Elements	48
Errors	49
See Also	50
ExportLicenseRecommendations	51
Request Syntax	51
Request Parameters	51
Response Syntax	54
Response Elements	54
Errors	55
See Also	56
ExportRDSDatabaseRecommendations	57
Request Syntax	57
Request Parameters	57
Response Syntax	62
Response Elements	62
Errors	63

See Also	64
GetAutoScalingGroupRecommendations	65
Request Syntax	65
Request Parameters	65
Response Syntax	67
Response Elements	69
Errors	70
See Also	71
GetEBSVolumeRecommendations	73
Request Syntax	73
Request Parameters	73
Response Syntax	74
Response Elements	76
Errors	77
See Also	78
GetEC2InstanceRecommendations	79
Request Syntax	79
Request Parameters	79
Response Syntax	81
Response Elements	83
Errors	84
See Also	85
GetEC2RecommendationProjectedMetrics	87
Request Syntax	87
Request Parameters	87
Response Syntax	88
Response Elements	89
Errors	89
See Also	90
GetECSServiceRecommendationProjectedMetrics	92
Request Syntax	92
Request Parameters	92
Response Syntax	93
Response Elements	93
Errors	94
See Also	95

GetECSServiceRecommendations	96
Request Syntax	96
Request Parameters	96
Response Syntax	97
Response Elements	100
Errors	100
See Also	101
GetEffectiveRecommendationPreferences	103
Request Syntax	103
Request Parameters	103
Response Syntax	103
Response Elements	104
Errors	106
See Also	107
GetEnrollmentStatus	108
Response Syntax	108
Response Elements	108
Errors	109
See Also	110
GetEnrollmentStatusesForOrganization	111
Request Syntax	111
Request Parameters	111
Response Syntax	112
Response Elements	112
Errors	113
See Also	113
GetIdleRecommendations	115
Request Syntax	115
Request Parameters	115
Response Syntax	117
Response Elements	118
Errors	118
See Also	119
GetLambdaFunctionRecommendations	121
Request Syntax	121
Request Parameters	121

Response Syntax	123
Response Elements	124
Errors	125
See Also	126
GetLicenseRecommendations	127
Request Syntax	127
Request Parameters	127
Response Syntax	128
Response Elements	130
Errors	130
See Also	131
GetRDSDatabaseRecommendationProjectedMetrics	133
Request Syntax	133
Request Parameters	133
Response Syntax	134
Response Elements	135
Errors	135
See Also	136
GetRDSDatabaseRecommendations	138
Request Syntax	138
Request Parameters	138
Response Syntax	140
Response Elements	143
Errors	143
See Also	144
GetRecommendationPreferences	146
Request Syntax	146
Request Parameters	146
Response Syntax	147
Response Elements	148
Errors	149
See Also	150
GetRecommendationSummaries	151
Request Syntax	151
Request Parameters	151
Response Syntax	152

Response Elements	154
Errors	154
See Also	155
PutRecommendationPreferences	156
Request Syntax	156
Request Parameters	157
Response Elements	161
Errors	161
See Also	162
UpdateEnrollmentStatus	163
Request Syntax	163
Request Parameters	163
Response Syntax	164
Response Elements	164
Errors	165
See Also	166
Data Types	167
AccountEnrollmentStatus	171
Contents	171
See Also	172
AutoScalingGroupConfiguration	173
Contents	173
See Also	174
AutoScalingGroupEstimatedMonthlySavings	175
Contents	175
See Also	175
AutoScalingGroupRecommendation	176
Contents	176
See Also	179
AutoScalingGroupRecommendationOption	180
Contents	180
See Also	182
AutoScalingGroupSavingsOpportunityAfterDiscounts	183
Contents	183
See Also	183
ContainerConfiguration	184

Contents	184
See Also	184
ContainerRecommendation	185
Contents	185
See Also	185
CurrentPerformanceRiskRatings	186
Contents	186
See Also	186
CustomizableMetricParameters	188
Contents	188
See Also	188
DBStorageConfiguration	190
Contents	190
See Also	191
EBSEffectiveRecommendationPreferences	192
Contents	192
See Also	192
EBSEstimatedMonthlySavings	193
Contents	193
See Also	193
EBSFilter	194
Contents	194
See Also	195
EBSSavingsEstimationMode	196
Contents	196
See Also	196
EBSSavingsOpportunityAfterDiscounts	197
Contents	197
See Also	197
EBSUtilizationMetric	198
Contents	198
See Also	199
ECSEffectiveRecommendationPreferences	200
Contents	200
See Also	200
ECSEstimatedMonthlySavings	201

Contents	201
See Also	201
ECSSavingsEstimationMode	202
Contents	202
See Also	202
ECSSavingsOpportunityAfterDiscounts	203
Contents	203
See Also	203
ECSServiceProjectedMetric	204
Contents	204
See Also	205
ECSServiceProjectedUtilizationMetric	206
Contents	206
See Also	207
ECSServiceRecommendation	208
Contents	208
See Also	211
ECSServiceRecommendationFilter	212
Contents	212
See Also	213
ECSServiceRecommendationOption	214
Contents	214
See Also	215
ECSServiceRecommendedOptionProjectedMetric	217
Contents	217
See Also	217
ECSServiceUtilizationMetric	219
Contents	219
See Also	220
EffectivePreferredResource	221
Contents	221
See Also	222
EffectiveRecommendationPreferences	223
Contents	223
See Also	225
EnrollmentFilter	226

Contents	226
See Also	226
EstimatedMonthlySavings	227
Contents	227
See Also	227
ExportDestination	228
Contents	228
See Also	228
ExternalMetricsPreference	229
Contents	229
See Also	229
ExternalMetricStatus	230
Contents	230
See Also	230
Filter	231
Contents	231
See Also	234
GetRecommendationError	235
Contents	235
See Also	235
Gpu	237
Contents	237
See Also	237
GpuInfo	238
Contents	238
See Also	238
IdleEstimatedMonthlySavings	239
Contents	239
See Also	239
IdleRecommendation	240
Contents	240
See Also	242
IdleRecommendationError	243
Contents	243
See Also	243
IdleRecommendationFilter	245

Contents	245
See Also	245
IdleSavingsOpportunity	247
Contents	247
See Also	247
IdleSavingsOpportunityAfterDiscounts	248
Contents	248
See Also	248
IdleSummary	249
Contents	249
See Also	249
IdleUtilizationMetric	250
Contents	250
See Also	251
InferredWorkloadSaving	252
Contents	252
See Also	253
InstanceEstimatedMonthlySavings	254
Contents	254
See Also	254
InstanceRecommendation	255
Contents	255
See Also	262
InstanceRecommendationOption	263
Contents	263
See Also	267
InstanceSavingsEstimationMode	268
Contents	268
See Also	268
InstanceSavingsOpportunityAfterDiscounts	269
Contents	269
See Also	269
JobFilter	270
Contents	270
See Also	271
LambdaEffectiveRecommendationPreferences	272

Contents	272
See Also	272
LambdaEstimatedMonthlySavings	273
Contents	273
See Also	273
LambdaFunctionMemoryProjectedMetric	274
Contents	274
See Also	274
LambdaFunctionMemoryRecommendationOption	276
Contents	276
See Also	277
LambdaFunctionRecommendation	278
Contents	278
See Also	282
LambdaFunctionRecommendationFilter	283
Contents	283
See Also	284
LambdaFunctionUtilizationMetric	285
Contents	285
See Also	286
LambdaSavingsEstimationMode	287
Contents	287
See Also	287
LambdaSavingsOpportunityAfterDiscounts	288
Contents	288
See Also	288
LicenseConfiguration	289
Contents	289
See Also	290
LicenseRecommendation	291
Contents	291
See Also	293
LicenseRecommendationFilter	294
Contents	294
See Also	295
LicenseRecommendationOption	296

Contents	296
See Also	297
MemorySizeConfiguration	298
Contents	298
See Also	298
MetricSource	299
Contents	299
See Also	299
OrderBy	300
Contents	300
See Also	300
PreferredResource	301
Contents	301
See Also	302
ProjectedMetric	303
Contents	303
See Also	305
RDSDatabaseProjectedMetric	306
Contents	306
See Also	307
RDSDatabaseRecommendedOptionProjectedMetric	308
Contents	308
See Also	308
RDSDBInstanceRecommendationOption	310
Contents	310
See Also	311
RDSDBRecommendation	313
Contents	313
See Also	317
RDSDBRecommendationFilter	318
Contents	318
See Also	319
RDSDBStorageRecommendationOption	320
Contents	320
See Also	321
RDSDBUtilizationMetric	322

Contents	322
See Also	323
RDSEffectiveRecommendationPreferences	324
Contents	324
See Also	325
RDSInstanceEstimatedMonthlySavings	326
Contents	326
See Also	326
RDSInstanceSavingsOpportunityAfterDiscounts	327
Contents	327
See Also	327
RDSSavingsEstimationMode	328
Contents	328
See Also	328
RDSSStorageEstimatedMonthlySavings	329
Contents	329
See Also	329
RDSSStorageSavingsOpportunityAfterDiscounts	330
Contents	330
See Also	330
ReasonCodeSummary	331
Contents	331
See Also	331
RecommendationExportJob	332
Contents	332
See Also	333
RecommendationPreferences	334
Contents	334
See Also	334
RecommendationPreferencesDetail	336
Contents	336
See Also	338
RecommendationSource	340
Contents	340
See Also	340
RecommendationSummary	341

Contents	341
See Also	343
RecommendedOptionProjectedMetric	344
Contents	344
See Also	345
S3Destination	346
Contents	346
See Also	346
S3DestinationConfig	348
Contents	348
See Also	348
SavingsOpportunity	349
Contents	349
See Also	350
Scope	351
Contents	351
See Also	352
ServiceConfiguration	353
Contents	353
See Also	354
Summary	355
Contents	355
See Also	355
Tag	356
Contents	356
See Also	356
UtilizationMetric	357
Contents	357
See Also	360
UtilizationPreference	361
Contents	361
See Also	361
VolumeConfiguration	362
Contents	362
See Also	363
VolumeRecommendation	364

- Contents 364
- See Also 366
- VolumeRecommendationOption 367
 - Contents 367
 - See Also 368
- Common Parameters 369**
- Common Errors 372**

Welcome

AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of your AWS compute resources, such as Amazon EC2 instances, Amazon EC2 Auto Scaling groups, AWS Lambda functions, Amazon EBS volumes, and Amazon ECS services on Fargate. It reports whether your resources are optimal, and generates optimization recommendations to reduce the cost and improve the performance of your workloads. Compute Optimizer also provides recent utilization metric data, in addition to projected utilization metric data for the recommendations, which you can use to evaluate which recommendation provides the best price-performance trade-off. The analysis of your usage patterns can help you decide when to move or resize your running resources, and still meet your performance and capacity requirements. For more information about Compute Optimizer, including the required permissions to use the service, see the [AWS Compute Optimizer User Guide](#).

This document was last published on September 3, 2025.

Actions

The following actions are supported:

- [DeleteRecommendationPreferences](#)
- [DescribeRecommendationExportJobs](#)
- [ExportAutoScalingGroupRecommendations](#)
- [ExportEBSVolumeRecommendations](#)
- [ExportEC2InstanceRecommendations](#)
- [ExportECSServiceRecommendations](#)
- [ExportIdleRecommendations](#)
- [ExportLambdaFunctionRecommendations](#)
- [ExportLicenseRecommendations](#)
- [ExportRDSDatabaseRecommendations](#)
- [GetAutoScalingGroupRecommendations](#)
- [GetEBSVolumeRecommendations](#)
- [GetEC2InstanceRecommendations](#)
- [GetEC2RecommendationProjectedMetrics](#)
- [GetECSServiceRecommendationProjectedMetrics](#)
- [GetECSServiceRecommendations](#)
- [GetEffectiveRecommendationPreferences](#)
- [GetEnrollmentStatus](#)
- [GetEnrollmentStatusesForOrganization](#)
- [GetIdleRecommendations](#)
- [GetLambdaFunctionRecommendations](#)
- [GetLicenseRecommendations](#)
- [GetRDSDatabaseRecommendationProjectedMetrics](#)
- [GetRDSDatabaseRecommendations](#)
- [GetRecommendationPreferences](#)
- [GetRecommendationSummaries](#)
- [PutRecommendationPreferences](#)

- [UpdateEnrollmentStatus](#)

DeleteRecommendationPreferences

Deletes a recommendation preference, such as enhanced infrastructure metrics.

For more information, see [Activating enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Request Syntax

```
{
  "recommendationPreferenceNames": [ "string" ],
  "resourceType": "string",
  "scope": {
    "name": "string",
    "value": "string"
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[recommendationPreferenceNames](#)

The name of the recommendation preference to delete.

Type: Array of strings

Valid Values: EnhancedInfrastructureMetrics | InferredWorkloadTypes
| ExternalMetricsPreference | LookBackPeriodPreference |
PreferredResources | UtilizationPreferences

Required: Yes

[resourceType](#)

The target resource type of the recommendation preference to delete.

The Ec2Instance option encompasses standalone instances and instances that are part of Auto Scaling groups. The AutoScalingGroup option encompasses only instances that are part of an Auto Scaling group.

Type: String

Valid Values: Ec2Instance | AutoScalingGroup | RdsDBInstance | AuroraDBClusterStorage

Required: Yes

scope

An object that describes the scope of the recommendation preference to delete.

You can delete recommendation preferences that are created at the organization level (for management accounts of an organization only), account level, and resource level. For more information, see [Activating enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Type: [Scope](#) object

Required: No

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

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

DescribeRecommendationExportJobs

Describes recommendation export jobs created in the last seven days.

Use the [ExportAutoScalingGroupRecommendations](#) or [ExportEC2InstanceRecommendations](#) actions to request an export of your recommendations. Then use the [DescribeRecommendationExportJobs](#) action to view your export jobs.

Request Syntax

```
{
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "jobIds": [ "string" ],
  "maxResults": number,
  "nextToken": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

filters

An array of objects to specify a filter that returns a more specific list of export jobs.

Type: Array of [JobFilter](#) objects

Required: No

jobIds

The identification numbers of the export jobs to return.

An export job ID is returned when you create an export using the [ExportAutoScalingGroupRecommendations](#) or [ExportEC2InstanceRecommendations](#) actions.

All export jobs created in the last seven days are returned if this parameter is omitted.

Type: Array of strings

Required: No

maxResults

The maximum number of export jobs to return with a single request.

To retrieve the remaining results, make another request with the returned nextToken value.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

nextToken

The token to advance to the next page of export jobs.

Type: String

Required: No

Response Syntax

```
{
  "nextToken": "string",
  "recommendationExportJobs": [
    {
      "creationTimestamp": number,
      "destination": {
        "s3": {
          "bucket": "string",
          "key": "string",
          "metadataKey": "string"
        }
      },
      "failureReason": "string",
      "jobId": "string",
      "lastUpdatedTimestamp": number,
      "resourceType": "string",
    }
  ]
}
```

```
    "status": "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.

nextToken

The token to use to advance to the next page of export jobs.

This value is null when there are no more pages of export jobs to return.

Type: String

recommendationExportJobs

An array of objects that describe recommendation export jobs.

Type: Array of [RecommendationExportJob](#) objects

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerError

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

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

ExportAutoScalingGroupRecommendations

Exports optimization recommendations for Auto Scaling groups.

Recommendations are exported in a comma-separated values (.csv) file, and its metadata in a JavaScript Object Notation (JSON) (.json) file, to an existing Amazon Simple Storage Service (Amazon S3) bucket that you specify. For more information, see [Exporting Recommendations](#) in the *Compute Optimizer User Guide*.

You can have only one Auto Scaling group export job in progress per AWS Region.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "fieldsToExport": [ "string" ],
  "fileFormat": "string",
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "includeMemberAccounts": boolean,
  "recommendationPreferences": {
    "cpuVendorArchitectures": [ "string" ]
  },
  "s3DestinationConfig": {
    "bucket": "string",
    "keyPrefix": "string"
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

accountIds

The IDs of the AWS accounts for which to export Auto Scaling group recommendations.

If your account is the management account of an organization, use this parameter to specify the member account for which you want to export recommendations.

This parameter cannot be specified together with the `include member accounts` parameter. The parameters are mutually exclusive.

Recommendations for member accounts are not included in the export if this parameter, or the `include member accounts` parameter, is omitted.

You can specify multiple account IDs per request.

Type: Array of strings

Required: No

fieldsToExport

The recommendations data to include in the export file. For more information about the fields that can be exported, see [Exported files](#) in the *Compute Optimizer User Guide*.

Type: Array of strings

Valid Values: `AccountId` | `AutoScalingGroupArn` | `AutoScalingGroupName` | `Finding` | `UtilizationMetricsCpuMaximum` | `UtilizationMetricsMemoryMaximum` | `UtilizationMetricsEbsReadOpsPerSecondMaximum` | `UtilizationMetricsEbsWriteOpsPerSecondMaximum` | `UtilizationMetricsEbsReadBytesPerSecondMaximum` | `UtilizationMetricsEbsWriteBytesPerSecondMaximum` | `UtilizationMetricsDiskReadOpsPerSecondMaximum` | `UtilizationMetricsDiskWriteOpsPerSecondMaximum` | `UtilizationMetricsDiskReadBytesPerSecondMaximum` | `UtilizationMetricsDiskWriteBytesPerSecondMaximum` | `UtilizationMetricsNetworkInBytesPerSecondMaximum` | `UtilizationMetricsNetworkOutBytesPerSecondMaximum` | `UtilizationMetricsNetworkPacketsInPerSecondMaximum` | `UtilizationMetricsNetworkPacketsOutPerSecondMaximum` | `LookbackPeriodInDays` | `CurrentConfigurationInstanceType` | `CurrentConfigurationDesiredCapacity` | `CurrentConfigurationMinSize` | `CurrentConfigurationMaxSize` | `CurrentConfigurationAllocationStrategy` | `CurrentConfigurationMixedInstanceTypes` | `CurrentConfigurationType` | `CurrentOnDemandPrice` | `CurrentStandardOneYearNoUpfrontReservedPrice`

| CurrentStandardThreeYearNoUpfrontReservedPrice |
CurrentVCpus | CurrentMemory | CurrentStorage | CurrentNetwork
| RecommendationOptionsConfigurationInstanceType |
RecommendationOptionsConfigurationDesiredCapacity
| RecommendationOptionsConfigurationMinSize |
RecommendationOptionsConfigurationMaxSize |
RecommendationOptionsConfigurationEstimatedInstanceHourReductionPercentage
| RecommendationOptionsConfigurationAllocationStrategy
| RecommendationOptionsConfigurationMixedInstanceTypes |
RecommendationOptionsConfigurationType |
RecommendationOptionsProjectedUtilizationMetricsCpuMaximum |
RecommendationOptionsProjectedUtilizationMetricsMemoryMaximum |
RecommendationOptionsPerformanceRisk |
RecommendationOptionsOnDemandPrice |
RecommendationOptionsStandardOneYearNoUpfrontReservedPrice |
RecommendationOptionsStandardThreeYearNoUpfrontReservedPrice
| RecommendationOptionsVcpus | RecommendationOptionsMemory |
RecommendationOptionsStorage | RecommendationOptionsNetwork
| LastRefreshTimestamp | CurrentPerformanceRisk |
RecommendationOptionsSavingsOpportunityPercentage |
RecommendationOptionsEstimatedMonthlySavingsCurrency
| RecommendationOptionsEstimatedMonthlySavingsValue |
EffectiveRecommendationPreferencesCpuVendorArchitectures |
EffectiveRecommendationPreferencesEnhancedInfrastructureMetrics
| EffectiveRecommendationPreferencesInferredWorkloadTypes
| EffectiveRecommendationPreferencesPreferredResources
| EffectiveRecommendationPreferencesLookBackPeriod |
InferredWorkloadTypes | RecommendationOptionsMigrationEffort |
CurrentInstanceGpuInfo | RecommendationOptionsInstanceGpuInfo |
UtilizationMetricsGpuPercentageMaximum |
UtilizationMetricsGpuMemoryPercentageMaximum |
RecommendationOptionsProjectedUtilizationMetricsGpuPercentageMaximum |
RecommendationOptionsProjectedUtilizationMetricsGpuMemoryPercentageMaximum
| EffectiveRecommendationPreferencesSavingsEstimationMode |
RecommendationOptionsSavingsOpportunityAfterDiscountsPercentage |

RecommendationOptionsEstimatedMonthlySavingsCurrencyAfterDiscounts |
RecommendationOptionsEstimatedMonthlySavingsValueAfterDiscounts

Required: No

fileFormat

The format of the export file.

The only export file format currently supported is Csv.

Type: String

Valid Values: Csv

Required: No

filters

An array of objects to specify a filter that exports a more specific set of Auto Scaling group recommendations.

Type: Array of [Filter](#) objects

Required: No

includeMemberAccounts

Indicates whether to include recommendations for resources in all member accounts of the organization if your account is the management account of an organization.

The member accounts must also be opted in to Compute Optimizer, and trusted access for Compute Optimizer must be enabled in the organization account. For more information, see [Compute Optimizer and AWS Organizations trusted access](#) in the *AWS Compute Optimizer User Guide*.

Recommendations for member accounts of the organization are not included in the export file if this parameter is omitted.

This parameter cannot be specified together with the account IDs parameter. The parameters are mutually exclusive.

Recommendations for member accounts are not included in the export if this parameter, or the account IDs parameter, is omitted.

Type: Boolean

Required: No

recommendationPreferences

An object to specify the preferences for the Auto Scaling group recommendations to export.

Type: [RecommendationPreferences](#) object

Required: No

s3DestinationConfig

An object to specify the destination Amazon Simple Storage Service (Amazon S3) bucket name and key prefix for the export job.

You must create the destination Amazon S3 bucket for your recommendations export before you create the export job. Compute Optimizer does not create the S3 bucket for you. After you create the S3 bucket, ensure that it has the required permissions policy to allow Compute Optimizer to write the export file to it. If you plan to specify an object prefix when you create the export job, you must include the object prefix in the policy that you add to the S3 bucket. For more information, see [Amazon S3 Bucket Policy for Compute Optimizer](#) in the *Compute Optimizer User Guide*.

Type: [S3DestinationConfig](#) object

Required: Yes

Response Syntax

```
{
  "jobId": "string",
  "s3Destination": {
    "bucket": "string",
    "key": "string",
    "metadataKey": "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

The identification number of the export job.

Use the [DescribeRecommendationExportJobs](#) action, and specify the job ID to view the status of an export job.

Type: String

s3Destination

An object that describes the destination Amazon S3 bucket of a recommendations export file.

Type: [S3Destination](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

LimitExceededException

The request exceeds a limit of the service.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ExportEBSVolumeRecommendations

Exports optimization recommendations for Amazon EBS volumes.

Recommendations are exported in a comma-separated values (.csv) file, and its metadata in a JavaScript Object Notation (JSON) (.json) file, to an existing Amazon Simple Storage Service (Amazon S3) bucket that you specify. For more information, see [Exporting Recommendations](#) in the *Compute Optimizer User Guide*.

You can have only one Amazon EBS volume export job in progress per AWS Region.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "fieldsToExport": [ "string" ],
  "fileFormat": "string",
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "includeMemberAccounts": boolean,
  "s3DestinationConfig": {
    "bucket": "string",
    "keyPrefix": "string"
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

accountIds

The IDs of the AWS accounts for which to export Amazon EBS volume recommendations.

If your account is the management account of an organization, use this parameter to specify the member account for which you want to export recommendations.

This parameter cannot be specified together with the `includeMemberAccounts` parameter. The parameters are mutually exclusive.

Recommendations for member accounts are not included in the export if this parameter, or the `includeMemberAccounts` parameter, is omitted.

You can specify multiple account IDs per request.

Type: Array of strings

Required: No

fieldsToExport

The recommendations data to include in the export file. For more information about the fields that can be exported, see [Exported files](#) in the *Compute Optimizer User Guide*.

Type: Array of strings

Valid Values: `AccountId` | `VolumeArn` | `Finding` | `UtilizationMetricsVolumeReadOpsPerSecondMaximum` | `UtilizationMetricsVolumeWriteOpsPerSecondMaximum` | `UtilizationMetricsVolumeReadBytesPerSecondMaximum` | `UtilizationMetricsVolumeWriteBytesPerSecondMaximum` | `LookbackPeriodInDays` | `CurrentConfigurationVolumeType` | `CurrentConfigurationVolumeBaselineIOPS` | `CurrentConfigurationVolumeBaselineThroughput` | `CurrentConfigurationVolumeBurstIOPS` | `CurrentConfigurationVolumeBurstThroughput` | `CurrentConfigurationVolumeSize` | `CurrentMonthlyPrice` | `RecommendationOptionsConfigurationVolumeType` | `RecommendationOptionsConfigurationVolumeBaselineIOPS` | `RecommendationOptionsConfigurationVolumeBaselineThroughput` | `RecommendationOptionsConfigurationVolumeBurstIOPS` | `RecommendationOptionsConfigurationVolumeBurstThroughput` | `RecommendationOptionsConfigurationVolumeSize` | `RecommendationOptionsMonthlyPrice` | `RecommendationOptionsPerformanceRisk` | `LastRefreshTimestamp` | `CurrentPerformanceRisk` | `RecommendationOptionsSavingsOpportunityPercentage` | `RecommendationOptionsEstimatedMonthlySavingsCurrency`

| RecommendationOptionsEstimatedMonthlySavingsValue |
Tags | RootVolume | CurrentConfigurationRootVolume |
EffectiveRecommendationPreferencesSavingsEstimationMode |
RecommendationOptionsSavingsOpportunityAfterDiscountsPercentage |
RecommendationOptionsEstimatedMonthlySavingsCurrencyAfterDiscounts |
RecommendationOptionsEstimatedMonthlySavingsValueAfterDiscounts

Required: No

fileFormat

The format of the export file.

The only export file format currently supported is Csv.

Type: String

Valid Values: Csv

Required: No

filters

An array of objects to specify a filter that exports a more specific set of Amazon EBS volume recommendations.

Type: Array of [EBSFilter](#) objects

Required: No

includeMemberAccounts

Indicates whether to include recommendations for resources in all member accounts of the organization if your account is the management account of an organization.

The member accounts must also be opted in to Compute Optimizer, and trusted access for Compute Optimizer must be enabled in the organization account. For more information, see [Compute Optimizer and AWS Organizations trusted access](#) in the *AWS Compute Optimizer User Guide*.

Recommendations for member accounts of the organization are not included in the export file if this parameter is omitted.

This parameter cannot be specified together with the account IDs parameter. The parameters are mutually exclusive.

Recommendations for member accounts are not included in the export if this parameter, or the account IDs parameter, is omitted.

Type: Boolean

Required: No

s3DestinationConfig

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and key prefix for a recommendations export job.

You must create the destination Amazon S3 bucket for your recommendations export before you create the export job. Compute Optimizer does not create the S3 bucket for you. After you create the S3 bucket, ensure that it has the required permission policy to allow Compute Optimizer to write the export file to it. If you plan to specify an object prefix when you create the export job, you must include the object prefix in the policy that you add to the S3 bucket. For more information, see [Amazon S3 Bucket Policy for Compute Optimizer](#) in the *Compute Optimizer User Guide*.

Type: [S3DestinationConfig](#) object

Required: Yes

Response Syntax

```
{
  "jobId": "string",
  "s3Destination": {
    "bucket": "string",
    "key": "string",
    "metadataKey": "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

The identification number of the export job.

Use the [DescribeRecommendationExportJobs](#) action, and specify the job ID to view the status of an export job.

Type: String

s3Destination

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and object keys of a recommendations export file, and its associated metadata file.

Type: [S3Destination](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

LimitExceededException

The request exceeds a limit of the service.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ExportEC2InstanceRecommendations

Exports optimization recommendations for Amazon EC2 instances.

Recommendations are exported in a comma-separated values (.csv) file, and its metadata in a JavaScript Object Notation (JSON) (.json) file, to an existing Amazon Simple Storage Service (Amazon S3) bucket that you specify. For more information, see [Exporting Recommendations](#) in the *Compute Optimizer User Guide*.

You can have only one Amazon EC2 instance export job in progress per AWS Region.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "fieldsToExport": [ "string" ],
  "fileFormat": "string",
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "includeMemberAccounts": boolean,
  "recommendationPreferences": {
    "cpuVendorArchitectures": [ "string" ]
  },
  "s3DestinationConfig": {
    "bucket": "string",
    "keyPrefix": "string"
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

accountIds

The IDs of the AWS accounts for which to export instance recommendations.

If your account is the management account of an organization, use this parameter to specify the member account for which you want to export recommendations.

This parameter cannot be specified together with the `include member accounts` parameter. The parameters are mutually exclusive.

Recommendations for member accounts are not included in the export if this parameter, or the `include member accounts` parameter, is omitted.

You can specify multiple account IDs per request.

Type: Array of strings

Required: No

fieldsToExport

The recommendations data to include in the export file. For more information about the fields that can be exported, see [Exported files](#) in the *Compute Optimizer User Guide*.

Type: Array of strings

Valid Values: `AccountId` | `InstanceArn` | `InstanceName` | `Finding` | `FindingReasonCodes` | `LookbackPeriodInDays` | `CurrentInstanceType` | `UtilizationMetricsCpuMaximum` | `UtilizationMetricsMemoryMaximum` | `UtilizationMetricsEbsReadOpsPerSecondMaximum` | `UtilizationMetricsEbsWriteOpsPerSecondMaximum` | `UtilizationMetricsEbsReadBytesPerSecondMaximum` | `UtilizationMetricsEbsWriteBytesPerSecondMaximum` | `UtilizationMetricsDiskReadOpsPerSecondMaximum` | `UtilizationMetricsDiskWriteOpsPerSecondMaximum` | `UtilizationMetricsDiskReadBytesPerSecondMaximum` | `UtilizationMetricsDiskWriteBytesPerSecondMaximum` | `UtilizationMetricsNetworkInBytesPerSecondMaximum` | `UtilizationMetricsNetworkOutBytesPerSecondMaximum` | `UtilizationMetricsNetworkPacketsInPerSecondMaximum` | `UtilizationMetricsNetworkPacketsOutPerSecondMaximum` | `CurrentOnDemandPrice` | `CurrentStandardOneYearNoUpfrontReservedPrice` | `CurrentStandardThreeYearNoUpfrontReservedPrice` | `CurrentVCpus` | `CurrentMemory` | `CurrentStorage` | `CurrentNetwork` | `RecommendationOptionsInstanceType` |

RecommendationOptionsProjectedUtilizationMetricsCpuMaximum |
 RecommendationOptionsProjectedUtilizationMetricsMemoryMaximum |
 RecommendationOptionsPlatformDifferences |
 RecommendationOptionsPerformanceRisk | RecommendationOptionsVcpus
 | RecommendationOptionsMemory | RecommendationOptionsStorage |
 RecommendationOptionsNetwork | RecommendationOptionsOnDemandPrice
 | RecommendationOptionsStandardOneYearNoUpfrontReservedPrice
 | RecommendationOptionsStandardThreeYearNoUpfrontReservedPrice
 | RecommendationsSourcesRecommendationSourceArn
 | RecommendationsSourcesRecommendationSourceType
 | LastRefreshTimestamp | CurrentPerformanceRisk |
 RecommendationOptionsSavingsOpportunityPercentage |
 RecommendationOptionsEstimatedMonthlySavingsCurrency
 | RecommendationOptionsEstimatedMonthlySavingsValue |
 EffectiveRecommendationPreferencesCpuVendorArchitectures |
 EffectiveRecommendationPreferencesEnhancedInfrastructureMetrics
 | EffectiveRecommendationPreferencesInferredWorkloadTypes |
 InferredWorkloadTypes | RecommendationOptionsMigrationEffort |
 EffectiveRecommendationPreferencesExternalMetricsSource | Tags |
 InstanceState | ExternalMetricStatusCode | ExternalMetricStatusReason
 | CurrentInstanceGpuInfo | RecommendationOptionsInstanceGpuInfo |
 UtilizationMetricsGpuPercentageMaximum |
 UtilizationMetricsGpuMemoryPercentageMaximum |
 RecommendationOptionsProjectedUtilizationMetricsGpuPercentageMaximum |
 RecommendationOptionsProjectedUtilizationMetricsGpuMemoryPercentageMaximum
 | Idle | EffectiveRecommendationPreferencesPreferredResources
 | EffectiveRecommendationPreferencesLookBackPeriod |
 EffectiveRecommendationPreferencesUtilizationPreferences |
 EffectiveRecommendationPreferencesSavingsEstimationMode |
 RecommendationOptionsSavingsOpportunityAfterDiscountsPercentage |
 RecommendationOptionsEstimatedMonthlySavingsCurrencyAfterDiscounts |
 RecommendationOptionsEstimatedMonthlySavingsValueAfterDiscounts

Required: No

fileFormat

The format of the export file.

The only export file format currently supported is Csv.

Type: String

Valid Values: Csv

Required: No

filters

An array of objects to specify a filter that exports a more specific set of instance recommendations.

Type: Array of [Filter](#) objects

Required: No

includeMemberAccounts

Indicates whether to include recommendations for resources in all member accounts of the organization if your account is the management account of an organization.

The member accounts must also be opted in to Compute Optimizer, and trusted access for Compute Optimizer must be enabled in the organization account. For more information, see [Compute Optimizer and AWS Organizations trusted access](#) in the *AWS Compute Optimizer User Guide*.

Recommendations for member accounts of the organization are not included in the export file if this parameter is omitted.

Recommendations for member accounts are not included in the export if this parameter, or the account IDs parameter, is omitted.

Type: Boolean

Required: No

recommendationPreferences

An object to specify the preferences for the Amazon EC2 instance recommendations to export.

Type: [RecommendationPreferences](#) object

Required: No

s3DestinationConfig

An object to specify the destination Amazon Simple Storage Service (Amazon S3) bucket name and key prefix for the export job.

You must create the destination Amazon S3 bucket for your recommendations export before you create the export job. Compute Optimizer does not create the S3 bucket for you. After you create the S3 bucket, ensure that it has the required permissions policy to allow Compute Optimizer to write the export file to it. If you plan to specify an object prefix when you create the export job, you must include the object prefix in the policy that you add to the S3 bucket. For more information, see [Amazon S3 Bucket Policy for Compute Optimizer](#) in the *Compute Optimizer User Guide*.

Type: [S3DestinationConfig](#) object

Required: Yes

Response Syntax

```
{
  "jobId": "string",
  "s3Destination": {
    "bucket": "string",
    "key": "string",
    "metadataKey": "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

The identification number of the export job.

Use the [DescribeRecommendationExportJobs](#) action, and specify the job ID to view the status of an export job.

Type: String

s3Destination

An object that describes the destination Amazon S3 bucket of a recommendations export file.

Type: [S3Destination](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

LimitExceededException

The request exceeds a limit of the service.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ExportECSServiceRecommendations

Exports optimization recommendations for Amazon ECS services on Fargate.

Recommendations are exported in a CSV file, and its metadata in a JSON file, to an existing Amazon Simple Storage Service (Amazon S3) bucket that you specify. For more information, see [Exporting Recommendations](#) in the *Compute Optimizer User Guide*.

You can only have one Amazon ECS service export job in progress per AWS Region.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "fieldsToExport": [ "string" ],
  "fileFormat": "string",
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "includeMemberAccounts": boolean,
  "s3DestinationConfig": {
    "bucket": "string",
    "keyPrefix": "string"
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[accountIds](#)

The AWS account IDs for the export Amazon ECS service recommendations.

If your account is the management account or the delegated administrator of an organization, use this parameter to specify the member account you want to export recommendations to.

This parameter can't be specified together with the `includeMemberAccounts` parameter. The parameters are mutually exclusive.

If this parameter or the `includeMemberAccounts` parameter is omitted, the recommendations for member accounts aren't included in the export.

You can specify multiple account IDs per request.

Type: Array of strings

Required: No

fieldsToExport

The recommendations data to include in the export file. For more information about the fields that can be exported, see [Exported files](#) in the *Compute Optimizer User Guide*.

Type: Array of strings

Valid Values: `AccountId` | `ServiceArn` | `LookbackPeriodInDays` | `LastRefreshTimestamp` | `LaunchType` | `CurrentPerformanceRisk` | `CurrentServiceConfigurationMemory` | `CurrentServiceConfigurationCpu` | `CurrentServiceConfigurationTaskDefinitionArn` | `CurrentServiceConfigurationAutoScalingConfiguration` | `CurrentServiceContainerConfigurations` | `UtilizationMetricsCpuMaximum` | `UtilizationMetricsMemoryMaximum` | `Finding` | `FindingReasonCodes` | `RecommendationOptionsMemory` | `RecommendationOptionsCpu` | `RecommendationOptionsSavingsOpportunityPercentage` | `RecommendationOptionsEstimatedMonthlySavingsCurrency` | `RecommendationOptionsEstimatedMonthlySavingsValue` | `RecommendationOptionsContainerRecommendations` | `RecommendationOptionsProjectedUtilizationMetricsCpuMaximum` | `RecommendationOptionsProjectedUtilizationMetricsMemoryMaximum` | `Tags` | `EffectiveRecommendationPreferencesSavingsEstimationMode` | `RecommendationOptionsSavingsOpportunityAfterDiscountsPercentage` | `RecommendationOptionsEstimatedMonthlySavingsCurrencyAfterDiscounts` | `RecommendationOptionsEstimatedMonthlySavingsValueAfterDiscounts`

Required: No

fileFormat

The format of the export file.

The CSV file is the only export file format currently supported.

Type: String

Valid Values: Csv

Required: No

filters

An array of objects to specify a filter that exports a more specific set of Amazon ECS service recommendations.

Type: Array of [ECSServiceRecommendationFilter](#) objects

Required: No

includeMemberAccounts

If your account is the management account or the delegated administrator of an organization, this parameter indicates whether to include recommendations for resources in all member accounts of the organization.

The member accounts must also be opted in to Compute Optimizer, and trusted access for Compute Optimizer must be enabled in the organization account. For more information, see [Compute Optimizer and AWS Organizations trusted access](#) in the *AWS Compute Optimizer User Guide*.

If this parameter is omitted, recommendations for member accounts of the organization aren't included in the export file.

If this parameter or the account ID parameter is omitted, recommendations for member accounts aren't included in the export.

Type: Boolean

Required: No

s3DestinationConfig

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and key prefix for a recommendations export job.

You must create the destination Amazon S3 bucket for your recommendations export before you create the export job. Compute Optimizer does not create the S3 bucket for you. After you create the S3 bucket, ensure that it has the required permission policy to allow Compute Optimizer to write the export file to it. If you plan to specify an object prefix when you create the export job, you must include the object prefix in the policy that you add to the S3 bucket. For more information, see [Amazon S3 Bucket Policy for Compute Optimizer](#) in the *Compute Optimizer User Guide*.

Type: [S3DestinationConfig](#) object

Required: Yes

Response Syntax

```
{
  "jobId": "string",
  "s3Destination": {
    "bucket": "string",
    "key": "string",
    "metadataKey": "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](#)

The identification number of the export job.

To view the status of an export job, use the [DescribeRecommendationExportJobs](#) action and specify the job ID.

Type: String

[s3Destination](#)

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and object keys of a recommendations export file, and its associated metadata file.

Type: [S3Destination](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

LimitExceededException

The request exceeds a limit of the service.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ExportIdleRecommendations

Export optimization recommendations for your idle resources.

Recommendations are exported in a comma-separated values (CSV) file, and its metadata in a JavaScript Object Notation (JSON) file, to an existing Amazon Simple Storage Service (Amazon S3) bucket that you specify. For more information, see [Exporting Recommendations](#) in the *Compute Optimizer User Guide*.

You can have only one idle resource export job in progress per AWS Region.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "fieldsToExport": [ "string" ],
  "fileFormat": "string",
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "includeMemberAccounts": boolean,
  "s3DestinationConfig": {
    "bucket": "string",
    "keyPrefix": "string"
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[accountIds](#)

The AWS account IDs for the export idle resource recommendations.

If your account is the management account or the delegated administrator of an organization, use this parameter to specify the member account you want to export recommendations to.

This parameter can't be specified together with the `includeMemberAccounts` parameter. The parameters are mutually exclusive.

If this parameter or the `includeMemberAccounts` parameter is omitted, the recommendations for member accounts aren't included in the export.

You can specify multiple account IDs per request.

Type: Array of strings

Required: No

fieldsToExport

The recommendations data to include in the export file. For more information about the fields that can be exported, see [Exported files](#) in the *Compute Optimizer User Guide*.

Type: Array of strings

Valid Values: `AccountId` | `ResourceArn` | `ResourceId` | `ResourceType` | `LastRefreshTimestamp` | `LookbackPeriodInDays` | `SavingsOpportunity` | `SavingsOpportunityAfterDiscount` | `UtilizationMetricsCpuMaximum` | `UtilizationMetricsMemoryMaximum` | `UtilizationMetricsNetworkOutBytesPerSecondMaximum` | `UtilizationMetricsNetworkInBytesPerSecondMaximum` | `UtilizationMetricsDatabaseConnectionsMaximum` | `UtilizationMetricsEBSVolumeReadIOPSMaximum` | `UtilizationMetricsEBSVolumeWriteIOPSMaximum` | `UtilizationMetricsVolumeReadOpsPerSecondMaximum` | `UtilizationMetricsVolumeWriteOpsPerSecondMaximum` | `Finding` | `FindingDescription` | `Tags`

Required: No

fileFormat

The format of the export file. The CSV file is the only export file format currently supported.

Type: String

Valid Values: `Csv`

Required: No

filters

An array of objects to specify a filter that exports a more specific set of idle resource recommendations.

Type: Array of [IdleRecommendationFilter](#) objects

Required: No

includeMemberAccounts

If your account is the management account or the delegated administrator of an organization, this parameter indicates whether to include recommendations for resources in all member accounts of the organization.

The member accounts must also be opted in to Compute Optimizer, and trusted access for Compute Optimizer must be enabled in the organization account. For more information, see [Compute Optimizer and AWS Organizations trusted access](#) in the *AWS Compute Optimizer User Guide*.

If this parameter is omitted, recommendations for member accounts of the organization aren't included in the export file.

If this parameter or the account ID parameter is omitted, recommendations for member accounts aren't included in the export.

Type: Boolean

Required: No

s3DestinationConfig

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and key prefix for a recommendations export job.

You must create the destination Amazon S3 bucket for your recommendations export before you create the export job. Compute Optimizer does not create the S3 bucket for you. After you create the S3 bucket, ensure that it has the required permission policy to allow Compute Optimizer to write the export file to it. If you plan to specify an object prefix when you create the export job, you must include the object prefix in the policy that you add to the S3 bucket. For more information, see [Amazon S3 Bucket Policy for Compute Optimizer](#) in the *Compute Optimizer User Guide*.

Type: [S3DestinationConfig](#) object

Required: Yes

Response Syntax

```
{
  "jobId": "string",
  "s3Destination": {
    "bucket": "string",
    "key": "string",
    "metadataKey": "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](#)

The identification number of the export job.

To view the status of an export job, use the [DescribeRecommendationExportJobs](#) action and specify the job ID.

Type: String

[s3Destination](#)

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and object keys of a recommendations export file, and its associated metadata file.

Type: [S3Destination](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

LimitExceededException

The request exceeds a limit of the service.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ExportLambdaFunctionRecommendations

Exports optimization recommendations for AWS Lambda functions.

Recommendations are exported in a comma-separated values (.csv) file, and its metadata in a JavaScript Object Notation (JSON) (.json) file, to an existing Amazon Simple Storage Service (Amazon S3) bucket that you specify. For more information, see [Exporting Recommendations](#) in the *Compute Optimizer User Guide*.

You can have only one Lambda function export job in progress per AWS Region.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "fieldsToExport": [ "string" ],
  "fileFormat": "string",
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "includeMemberAccounts": boolean,
  "s3DestinationConfig": {
    "bucket": "string",
    "keyPrefix": "string"
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

accountIds

The IDs of the AWS accounts for which to export Lambda function recommendations.

If your account is the management account of an organization, use this parameter to specify the member account for which you want to export recommendations.

This parameter cannot be specified together with the `includeMemberAccounts` parameter. The parameters are mutually exclusive.

Recommendations for member accounts are not included in the export if this parameter, or the `includeMemberAccounts` parameter, is omitted.

You can specify multiple account IDs per request.

Type: Array of strings

Required: No

fieldsToExport

The recommendations data to include in the export file. For more information about the fields that can be exported, see [Exported files](#) in the *Compute Optimizer User Guide*.

Type: Array of strings

Valid Values: `AccountId` | `FunctionArn` | `FunctionVersion` | `Finding` | `FindingReasonCodes` | `NumberOfInvocations` | `UtilizationMetricsDurationMaximum` | `UtilizationMetricsDurationAverage` | `UtilizationMetricsMemoryMaximum` | `UtilizationMetricsMemoryAverage` | `LookbackPeriodInDays` | `CurrentConfigurationMemorySize` | `CurrentConfigurationTimeout` | `CurrentCostTotal` | `CurrentCostAverage` | `RecommendationOptionsConfigurationMemorySize` | `RecommendationOptionsCostLow` | `RecommendationOptionsCostHigh` | `RecommendationOptionsProjectedUtilizationMetricsDurationLowerBound` | `RecommendationOptionsProjectedUtilizationMetricsDurationUpperBound` | `RecommendationOptionsProjectedUtilizationMetricsDurationExpected` | `LastRefreshTimestamp` | `CurrentPerformanceRisk` | `RecommendationOptionsSavingsOpportunityPercentage` | `RecommendationOptionsEstimatedMonthlySavingsCurrency` | `RecommendationOptionsEstimatedMonthlySavingsValue` | `Tags` | `EffectiveRecommendationPreferencesSavingsEstimationMode` | `RecommendationOptionsSavingsOpportunityAfterDiscountsPercentage` | `RecommendationOptionsEstimatedMonthlySavingsCurrencyAfterDiscounts` | `RecommendationOptionsEstimatedMonthlySavingsValueAfterDiscounts`

Required: No

fileFormat

The format of the export file.

The only export file format currently supported is Csv.

Type: String

Valid Values: Csv

Required: No

filters

An array of objects to specify a filter that exports a more specific set of Lambda function recommendations.

Type: Array of [LambdaFunctionRecommendationFilter](#) objects

Required: No

includeMemberAccounts

Indicates whether to include recommendations for resources in all member accounts of the organization if your account is the management account of an organization.

The member accounts must also be opted in to Compute Optimizer, and trusted access for Compute Optimizer must be enabled in the organization account. For more information, see [Compute Optimizer and AWS Organizations trusted access](#) in the *AWS Compute Optimizer User Guide*.

Recommendations for member accounts of the organization are not included in the export file if this parameter is omitted.

This parameter cannot be specified together with the account IDs parameter. The parameters are mutually exclusive.

Recommendations for member accounts are not included in the export if this parameter, or the account IDs parameter, is omitted.

Type: Boolean

Required: No

s3DestinationConfig

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and key prefix for a recommendations export job.

You must create the destination Amazon S3 bucket for your recommendations export before you create the export job. Compute Optimizer does not create the S3 bucket for you. After you create the S3 bucket, ensure that it has the required permission policy to allow Compute Optimizer to write the export file to it. If you plan to specify an object prefix when you create the export job, you must include the object prefix in the policy that you add to the S3 bucket. For more information, see [Amazon S3 Bucket Policy for Compute Optimizer](#) in the *Compute Optimizer User Guide*.

Type: [S3DestinationConfig](#) object

Required: Yes

Response Syntax

```
{
  "jobId": "string",
  "s3Destination": {
    "bucket": "string",
    "key": "string",
    "metadataKey": "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

The identification number of the export job.

Use the [DescribeRecommendationExportJobs](#) action, and specify the job ID to view the status of an export job.

Type: String

s3Destination

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and object keys of a recommendations export file, and its associated metadata file.

Type: [S3Destination](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

LimitExceededException

The request exceeds a limit of the service.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ExportLicenseRecommendations

Export optimization recommendations for your licenses.

Recommendations are exported in a comma-separated values (CSV) file, and its metadata in a JavaScript Object Notation (JSON) file, to an existing Amazon Simple Storage Service (Amazon S3) bucket that you specify. For more information, see [Exporting Recommendations](#) in the *Compute Optimizer User Guide*.

You can have only one license export job in progress per AWS Region.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "fieldsToExport": [ "string" ],
  "fileFormat": "string",
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "includeMemberAccounts": boolean,
  "s3DestinationConfig": {
    "bucket": "string",
    "keyPrefix": "string"
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[accountIds](#)

The IDs of the AWS accounts for which to export license recommendations.

If your account is the management account of an organization, use this parameter to specify the member account for which you want to export recommendations.

This parameter can't be specified together with the `includeMemberAccounts` parameter. The parameters are mutually exclusive.

If this parameter is omitted, recommendations for member accounts aren't included in the export.

You can specify multiple account IDs per request.

Type: Array of strings

Required: No

fieldsToExport

The recommendations data to include in the export file. For more information about the fields that can be exported, see [Exported files](#) in the *Compute Optimizer User Guide*.

Type: Array of strings

Valid Values: `AccountId` | `ResourceArn` | `LookbackPeriodInDays` | `LastRefreshTimestamp` | `Finding` | `FindingReasonCodes` | `CurrentLicenseConfigurationNumberOfCores` | `CurrentLicenseConfigurationInstanceType` | `CurrentLicenseConfigurationOperatingSystem` | `CurrentLicenseConfigurationLicenseName` | `CurrentLicenseConfigurationLicenseEdition` | `CurrentLicenseConfigurationLicenseModel` | `CurrentLicenseConfigurationLicenseVersion` | `CurrentLicenseConfigurationMetricsSource` | `RecommendationOptionsOperatingSystem` | `RecommendationOptionsLicenseEdition` | `RecommendationOptionsLicenseModel` | `RecommendationOptionsSavingsOpportunityPercentage` | `RecommendationOptionsEstimatedMonthlySavingsCurrency` | `RecommendationOptionsEstimatedMonthlySavingsValue` | `Tags`

Required: No

fileFormat

The format of the export file.

A CSV file is the only export format currently supported.

Type: String

Valid Values: Csv

Required: No

filters

An array of objects to specify a filter that exports a more specific set of license recommendations.

Type: Array of [LicenseRecommendationFilter](#) objects

Required: No

includeMemberAccounts

Indicates whether to include recommendations for resources in all member accounts of the organization if your account is the management account of an organization.

The member accounts must also be opted in to Compute Optimizer, and trusted access for Compute Optimizer must be enabled in the organization account. For more information, see [Compute Optimizer and AWS Organizations trusted access](#) in the *AWS Compute Optimizer User Guide*.

If this parameter is omitted, recommendations for member accounts of the organization aren't included in the export file .

This parameter cannot be specified together with the account IDs parameter. The parameters are mutually exclusive.

Type: Boolean

Required: No

s3DestinationConfig

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and key prefix for a recommendations export job.

You must create the destination Amazon S3 bucket for your recommendations export before you create the export job. Compute Optimizer does not create the S3 bucket for you. After you create the S3 bucket, ensure that it has the required permission policy to allow Compute

Optimizer to write the export file to it. If you plan to specify an object prefix when you create the export job, you must include the object prefix in the policy that you add to the S3 bucket. For more information, see [Amazon S3 Bucket Policy for Compute Optimizer](#) in the *Compute Optimizer User Guide*.

Type: [S3DestinationConfig](#) object

Required: Yes

Response Syntax

```
{
  "jobId": "string",
  "s3Destination": {
    "bucket": "string",
    "key": "string",
    "metadataKey": "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](#)

The identification number of the export job.

To view the status of an export job, use the [DescribeRecommendationExportJobs](#) action and specify the job ID.

Type: String

[s3Destination](#)

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and object keys of a recommendations export file, and its associated metadata file.

Type: [S3Destination](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

LimitExceededException

The request exceeds a limit of the service.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ExportRDSDatabaseRecommendations

Export optimization recommendations for your Amazon Aurora and Amazon Relational Database Service (Amazon RDS) databases.

Recommendations are exported in a comma-separated values (CSV) file, and its metadata in a JavaScript Object Notation (JSON) file, to an existing Amazon Simple Storage Service (Amazon S3) bucket that you specify. For more information, see [Exporting Recommendations](#) in the *Compute Optimizer User Guide*.

You can have only one Amazon Aurora or RDS export job in progress per AWS Region.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "fieldsToExport": [ "string" ],
  "fileFormat": "string",
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "includeMemberAccounts": boolean,
  "recommendationPreferences": {
    "cpuVendorArchitectures": [ "string" ]
  },
  "s3DestinationConfig": {
    "bucket": "string",
    "keyPrefix": "string"
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

accountId

The AWS account IDs for the export Amazon Aurora and RDS database recommendations.

If your account is the management account or the delegated administrator of an organization, use this parameter to specify the member account you want to export recommendations to.

This parameter can't be specified together with the include member accounts parameter. The parameters are mutually exclusive.

If this parameter or the include member accounts parameter is omitted, the recommendations for member accounts aren't included in the export.

You can specify multiple account IDs per request.

Type: Array of strings

Required: No

fieldsToExport

The recommendations data to include in the export file. For more information about the fields that can be exported, see [Exported files](#) in the *Compute Optimizer User Guide*.

Type: Array of strings

Valid Values: ResourceArn | AccountId | Engine | EngineVersion | Idle
| MultiAZDBInstance | ClusterWriter | CurrentDBInstanceClass |
CurrentStorageConfigurationStorageType |
CurrentStorageConfigurationAllocatedStorage |
CurrentStorageConfigurationMaxAllocatedStorage |
CurrentStorageConfigurationIOPS |
CurrentStorageConfigurationStorageThroughput |
CurrentStorageEstimatedMonthlyVolumeIOPsCostVariation |
CurrentInstanceOnDemandHourlyPrice | CurrentStorageOnDemandMonthlyPrice
| LookbackPeriodInDays |
CurrentStorageEstimatedClusterInstanceOnDemandMonthlyCost |
CurrentStorageEstimatedClusterStorageOnDemandMonthlyCost |
CurrentStorageEstimatedClusterStorageIOOnDemandMonthlyCost |
CurrentInstancePerformanceRisk | UtilizationMetricsCpuMaximum |
UtilizationMetricsMemoryMaximum |
UtilizationMetricsEBSVolumeStorageSpaceUtilizationMaximum

| UtilizationMetricsNetworkReceiveThroughputMaximum
| UtilizationMetricsNetworkTransmitThroughputMaximum
| UtilizationMetricsEBSVolumeReadIOPSMaximum |
UtilizationMetricsEBSVolumeWriteIOPSMaximum |
UtilizationMetricsEBSVolumeReadThroughputMaximum |
UtilizationMetricsEBSVolumeWriteThroughputMaximum
| UtilizationMetricsDatabaseConnectionsMaximum |
UtilizationMetricsStorageNetworkReceiveThroughputMaximum |
UtilizationMetricsStorageNetworkTransmitThroughputMaximum
| UtilizationMetricsAuroraMemoryHealthStateMaximum |
UtilizationMetricsAuroraMemoryNumDeclinedSqlTotalMaximum
| UtilizationMetricsAuroraMemoryNumKillConnTotalMaximum
| UtilizationMetricsAuroraMemoryNumKillQueryTotalMaximum
| UtilizationMetricsReadIOPSEphemeralStorageMaximum |
UtilizationMetricsWriteIOPSEphemeralStorageMaximum |
UtilizationMetricsVolumeBytesUsedAverage |
UtilizationMetricsVolumeReadIOPsAverage |
UtilizationMetricsVolumeWriteIOPsAverage | InstanceFinding |
InstanceFindingReasonCodes | StorageFinding | StorageFindingReasonCodes
| InstanceRecommendationOptionsDBInstanceClass |
InstanceRecommendationOptionsRank |
InstanceRecommendationOptionsPerformanceRisk |
InstanceRecommendationOptionsProjectedUtilizationMetricsCpuMaximum |
StorageRecommendationOptionsStorageType |
StorageRecommendationOptionsAllocatedStorage |
StorageRecommendationOptionsMaxAllocatedStorage |
StorageRecommendationOptionsIOPS |
StorageRecommendationOptionsStorageThroughput |
StorageRecommendationOptionsRank |
StorageRecommendationOptionsEstimatedMonthlyVolumeIOPsCostVariation
| InstanceRecommendationOptionsInstanceOnDemandHourlyPrice |
InstanceRecommendationOptionsSavingsOpportunityPercentage |
InstanceRecommendationOptionsEstimatedMonthlySavingsCurrency
| InstanceRecommendationOptionsEstimatedMonthlySavingsValue |
InstanceRecommendationOptionsSavingsOpportunityAfterDiscountsPercentage
|

```

InstanceRecommendationOptionsEstimatedMonthlySavingsCurrencyAfterDiscounts
|
InstanceRecommendationOptionsEstimatedMonthlySavingsValueAfterDiscounts
| StorageRecommendationOptionsOnDemandMonthlyPrice |
StorageRecommendationOptionsEstimatedClusterInstanceOnDemandMonthlyCost
| StorageRecommendationOptionsEstimatedClusterStorageOnDemandMonthlyCost
|
StorageRecommendationOptionsEstimatedClusterStorageIOOnDemandMonthlyCost
| StorageRecommendationOptionsSavingsOpportunityPercentage |
StorageRecommendationOptionsEstimatedMonthlySavingsCurrency
| StorageRecommendationOptionsEstimatedMonthlySavingsValue |
StorageRecommendationOptionsSavingsOpportunityAfterDiscountsPercentage |
StorageRecommendationOptionsEstimatedMonthlySavingsCurrencyAfterDiscounts
| StorageRecommendationOptionsEstimatedMonthlySavingsValueAfterDiscounts
| EffectiveRecommendationPreferencesCpuVendorArchitectures |
EffectiveRecommendationPreferencesEnhancedInfrastructureMetrics
| EffectiveRecommendationPreferencesLookBackPeriod |
EffectiveRecommendationPreferencesSavingsEstimationMode |
LastRefreshTimestamp | Tags | DBClusterIdentifier | PromotionTier

```

Required: No

fileFormat

The format of the export file.

The CSV file is the only export file format currently supported.

Type: String

Valid Values: Csv

Required: No

filters

An array of objects to specify a filter that exports a more specific set of Amazon Aurora and RDS recommendations.

Type: Array of [RDSDBRecommendationFilter](#) objects

Required: No

includeMemberAccounts

If your account is the management account or the delegated administrator of an organization, this parameter indicates whether to include recommendations for resources in all member accounts of the organization.

The member accounts must also be opted in to Compute Optimizer, and trusted access for Compute Optimizer must be enabled in the organization account. For more information, see [Compute Optimizer and AWS Organizations trusted access](#) in the *AWS Compute Optimizer User Guide*.

If this parameter is omitted, recommendations for member accounts of the organization aren't included in the export file.

If this parameter or the account ID parameter is omitted, recommendations for member accounts aren't included in the export.

Type: Boolean

Required: No

recommendationPreferences

Describes the recommendation preferences to return in the response of a [GetAutoScalingGroupRecommendations](#), [GetEC2InstanceRecommendations](#), [GetEC2RecommendationProjectedMetrics](#), [GetRDSDatabaseRecommendations](#), and [GetRDSDatabaseRecommendationProjectedMetrics](#) request.

Type: [RecommendationPreferences](#) object

Required: No

s3DestinationConfig

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and key prefix for a recommendations export job.

You must create the destination Amazon S3 bucket for your recommendations export before you create the export job. Compute Optimizer does not create the S3 bucket for you. After you create the S3 bucket, ensure that it has the required permission policy to allow Compute

Optimizer to write the export file to it. If you plan to specify an object prefix when you create the export job, you must include the object prefix in the policy that you add to the S3 bucket. For more information, see [Amazon S3 Bucket Policy for Compute Optimizer](#) in the *Compute Optimizer User Guide*.

Type: [S3DestinationConfig](#) object

Required: Yes

Response Syntax

```
{
  "jobId": "string",
  "s3Destination": {
    "bucket": "string",
    "key": "string",
    "metadataKey": "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](#)

The identification number of the export job.

To view the status of an export job, use the [DescribeRecommendationExportJobs](#) action and specify the job ID.

Type: String

[s3Destination](#)

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and object keys of a recommendations export file, and its associated metadata file.

Type: [S3Destination](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

LimitExceededException

The request exceeds a limit of the service.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetAutoScalingGroupRecommendations

Returns Auto Scaling group recommendations.

AWS Compute Optimizer generates recommendations for Amazon EC2 Auto Scaling groups that meet a specific set of requirements. For more information, see the [Supported resources and requirements](#) in the *AWS Compute Optimizer User Guide*.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "autoScalingGroupArns": [ "string" ],
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "maxResults": number,
  "nextToken": "string",
  "recommendationPreferences": {
    "cpuVendorArchitectures": [ "string" ]
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[accountIds](#)

The ID of the AWS account for which to return Auto Scaling group recommendations.

If your account is the management account of an organization, use this parameter to specify the member account for which you want to return Auto Scaling group recommendations.

Only one account ID can be specified per request.

Type: Array of strings

Required: No

[autoScalingGroupArns](#)

The Amazon Resource Name (ARN) of the Auto Scaling groups for which to return recommendations.

Type: Array of strings

Required: No

[filters](#)

An array of objects to specify a filter that returns a more specific list of Auto Scaling group recommendations.

Type: Array of [Filter](#) objects

Required: No

[maxResults](#)

The maximum number of Auto Scaling group recommendations to return with a single request.

To retrieve the remaining results, make another request with the returned `nextToken` value.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

[nextToken](#)

The token to advance to the next page of Auto Scaling group recommendations.

Type: String

Required: No

[recommendationPreferences](#)

An object to specify the preferences for the Auto Scaling group recommendations to return in the response.

Type: [RecommendationPreferences](#) object

Required: No

Response Syntax

```
{
  "autoScalingGroupRecommendations": [
    {
      "accountId": "string",
      "autoScalingGroupArn": "string",
      "autoScalingGroupName": "string",
      "currentConfiguration": {
        "allocationStrategy": "string",
        "desiredCapacity": number,
        "estimatedInstanceHourReductionPercentage": number,
        "instanceType": "string",
        "maxSize": number,
        "minSize": number,
        "mixedInstanceTypes": [ "string" ],
        "type": "string"
      },
      "currentInstanceGpuInfo": {
        "gpus": [
          {
            "gpuCount": number,
            "gpuMemorySizeInMiB": number
          }
        ]
      },
      "currentPerformanceRisk": "string",
      "effectiveRecommendationPreferences": {
        "cpuVendorArchitectures": [ "string" ],
        "enhancedInfrastructureMetrics": "string",
        "externalMetricsPreference": {
          "source": "string"
        },
      },
      "inferredWorkloadTypes": "string",
      "lookBackPeriod": "string",
      "preferredResources": [
        {
          "effectiveIncludeList": [ "string" ],
          "excludeList": [ "string" ],
          "includeList": [ "string" ],
          "name": "string"
        }
      ]
    },
  ],
}
```

```

    "savingsEstimationMode": {
      "source": "string"
    },
    "utilizationPreferences": [
      {
        "metricName": "string",
        "metricParameters": {
          "headroom": "string",
          "threshold": "string"
        }
      }
    ]
  },
  "finding": "string",
  "inferredWorkloadTypes": [ "string" ],
  "lastRefreshTimestamp": number,
  "lookBackPeriodInDays": number,
  "recommendationOptions": [
    {
      "configuration": {
        "allocationStrategy": "string",
        "desiredCapacity": number,
        "estimatedInstanceHourReductionPercentage": number,
        "instanceType": "string",
        "maxSize": number,
        "minSize": number,
        "mixedInstanceTypes": [ "string" ],
        "type": "string"
      },
      "instanceGpuInfo": {
        "gpus": [
          {
            "gpuCount": number,
            "gpuMemorySizeInMiB": number
          }
        ]
      },
      "migrationEffort": "string",
      "performanceRisk": number,
      "projectedUtilizationMetrics": [
        {
          "name": "string",
          "statistic": "string",
          "value": number
        }
      ]
    }
  ]
}

```

```

    }
  ],
  "rank": number,
  "savingsOpportunity": {
    "estimatedMonthlySavings": {
      "currency": "string",
      "value": number
    },
    "savingsOpportunityPercentage": number
  },
  "savingsOpportunityAfterDiscounts": {
    "estimatedMonthlySavings": {
      "currency": "string",
      "value": number
    },
    "savingsOpportunityPercentage": number
  }
},
"utilizationMetrics": [
  {
    "name": "string",
    "statistic": "string",
    "value": number
  }
]
},
],
"errors": [
  {
    "code": "string",
    "identifier": "string",
    "message": "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.

[autoScalingGroupRecommendations](#)

An array of objects that describe Auto Scaling group recommendations.

Type: Array of [AutoScalingGroupRecommendation](#) objects

[errors](#)

An array of objects that describe errors of the request.

For example, an error is returned if you request recommendations for an unsupported Auto Scaling group.

Type: Array of [GetRecommendationError](#) objects

[nextToken](#)

The token to use to advance to the next page of Auto Scaling group recommendations.

This value is null when there are no more pages of Auto Scaling group recommendations to return.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerError

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)

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

GetEBSVolumeRecommendations

Returns Amazon Elastic Block Store (Amazon EBS) volume recommendations.

AWS Compute Optimizer generates recommendations for Amazon EBS volumes that meet a specific set of requirements. For more information, see the [Supported resources and requirements](#) in the *AWS Compute Optimizer User Guide*.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "maxResults": number,
  "nextToken": "string",
  "volumeArns": [ "string" ]
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

accountIds

The ID of the AWS account for which to return volume recommendations.

If your account is the management account of an organization, use this parameter to specify the member account for which you want to return volume recommendations.

Only one account ID can be specified per request.

Type: Array of strings

Required: No

filters

An array of objects to specify a filter that returns a more specific list of volume recommendations.

Type: Array of [EBSFilter](#) objects

Required: No

maxResults

The maximum number of volume recommendations to return with a single request.

To retrieve the remaining results, make another request with the returned `nextToken` value.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

nextToken

The token to advance to the next page of volume recommendations.

Type: String

Required: No

volumeArns

The Amazon Resource Name (ARN) of the volumes for which to return recommendations.

Type: Array of strings

Required: No

Response Syntax

```
{
  "errors": [
    {
      "code": "string",
      "identifier": "string",
      "message": "string"
    }
  ]
}
```

```
],
"nextToken": "string",
"volumeRecommendations": [
  {
    "accountId": "string",
    "currentConfiguration": {
      "rootVolume": boolean,
      "volumeBaselineIOPS": number,
      "volumeBaselineThroughput": number,
      "volumeBurstIOPS": number,
      "volumeBurstThroughput": number,
      "volumeSize": number,
      "volumeType": "string"
    },
    "currentPerformanceRisk": "string",
    "effectiveRecommendationPreferences": {
      "savingsEstimationMode": {
        "source": "string"
      }
    },
    "finding": "string",
    "lastRefreshTimestamp": number,
    "lookBackPeriodInDays": number,
    "tags": [
      {
        "key": "string",
        "value": "string"
      }
    ],
    "utilizationMetrics": [
      {
        "name": "string",
        "statistic": "string",
        "value": number
      }
    ],
    "volumeArn": "string",
    "volumeRecommendationOptions": [
      {
        "configuration": {
          "rootVolume": boolean,
          "volumeBaselineIOPS": number,
          "volumeBaselineThroughput": number,
          "volumeBurstIOPS": number,
```

```

        "volumeBurstThroughput": number,
        "volumeSize": number,
        "volumeType": "string"
    },
    "performanceRisk": number,
    "rank": number,
    "savingsOpportunity": {
        "estimatedMonthlySavings": {
            "currency": "string",
            "value": number
        },
        "savingsOpportunityPercentage": number
    },
    "savingsOpportunityAfterDiscounts": {
        "estimatedMonthlySavings": {
            "currency": "string",
            "value": number
        },
        "savingsOpportunityPercentage": 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.

errors

An array of objects that describe errors of the request.

For example, an error is returned if you request recommendations for an unsupported volume.

Type: Array of [GetRecommendationError](#) objects

nextToken

The token to use to advance to the next page of volume recommendations.

This value is null when there are no more pages of volume recommendations to return.

Type: String

volumeRecommendations

An array of objects that describe volume recommendations.

Type: Array of [VolumeRecommendation](#) objects

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetEC2InstanceRecommendations

Returns Amazon EC2 instance recommendations.

AWS Compute Optimizer generates recommendations for Amazon Elastic Compute Cloud (Amazon EC2) instances that meet a specific set of requirements. For more information, see the [Supported resources and requirements](#) in the *AWS Compute Optimizer User Guide*.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "instanceArns": [ "string" ],
  "maxResults": number,
  "nextToken": "string",
  "recommendationPreferences": {
    "cpuVendorArchitectures": [ "string" ]
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[accountIds](#)

The ID of the AWS account for which to return instance recommendations.

If your account is the management account of an organization, use this parameter to specify the member account for which you want to return instance recommendations.

Only one account ID can be specified per request.

Type: Array of strings

Required: No

filters

An array of objects to specify a filter that returns a more specific list of instance recommendations.

Type: Array of [Filter](#) objects

Required: No

instanceArns

The Amazon Resource Name (ARN) of the instances for which to return recommendations.

Type: Array of strings

Required: No

maxResults

The maximum number of instance recommendations to return with a single request.

To retrieve the remaining results, make another request with the returned `nextToken` value.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

nextToken

The token to advance to the next page of instance recommendations.

Type: String

Required: No

recommendationPreferences

An object to specify the preferences for the Amazon EC2 instance recommendations to return in the response.

Type: [RecommendationPreferences](#) object

Required: No

Response Syntax

```
{
  "errors": [
    {
      "code": "string",
      "identifier": "string",
      "message": "string"
    }
  ],
  "instanceRecommendations": [
    {
      "accountId": "string",
      "currentInstanceGpuInfo": {
        "gpus": [
          {
            "gpuCount": number,
            "gpuMemorySizeInMiB": number
          }
        ]
      },
      "currentInstanceType": "string",
      "currentPerformanceRisk": "string",
      "effectiveRecommendationPreferences": {
        "cpuVendorArchitectures": [ "string" ],
        "enhancedInfrastructureMetrics": "string",
        "externalMetricsPreference": {
          "source": "string"
        },
      },
      "inferredWorkloadTypes": "string",
      "lookBackPeriod": "string",
      "preferredResources": [
        {
          "effectiveIncludeList": [ "string" ],
          "excludeList": [ "string" ],
          "includeList": [ "string" ],
          "name": "string"
        }
      ],
      "savingsEstimationMode": {
        "source": "string"
      },
      "utilizationPreferences": [
```

```

        {
            "metricName": "string",
            "metricParameters": {
                "headroom": "string",
                "threshold": "string"
            }
        }
    ],
    "externalMetricStatus": {
        "statusCode": "string",
        "statusReason": "string"
    },
    "finding": "string",
    "findingReasonCodes": [ "string" ],
    "idle": "string",
    "inferredWorkloadTypes": [ "string" ],
    "instanceArn": "string",
    "instanceName": "string",
    "instanceState": "string",
    "lastRefreshTimestamp": number,
    "lookBackPeriodInDays": number,
    "recommendationOptions": [
        {
            "instanceGpuInfo": {
                "gpus": [
                    {
                        "gpuCount": number,
                        "gpuMemorySizeInMiB": number
                    }
                ]
            },
            "instanceType": "string",
            "migrationEffort": "string",
            "performanceRisk": number,
            "platformDifferences": [ "string" ],
            "projectedUtilizationMetrics": [
                {
                    "name": "string",
                    "statistic": "string",
                    "value": number
                }
            ],
            "rank": number,

```

```

    "savingsOpportunity": {
      "estimatedMonthlySavings": {
        "currency": "string",
        "value": number
      },
      "savingsOpportunityPercentage": number
    },
    "savingsOpportunityAfterDiscounts": {
      "estimatedMonthlySavings": {
        "currency": "string",
        "value": number
      },
      "savingsOpportunityPercentage": number
    }
  ],
  "recommendationSources": [
    {
      "recommendationSourceArn": "string",
      "recommendationSourceType": "string"
    }
  ],
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ],
  "utilizationMetrics": [
    {
      "name": "string",
      "statistic": "string",
      "value": 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.

errors

An array of objects that describe errors of the request.

For example, an error is returned if you request recommendations for an instance of an unsupported instance family.

Type: Array of [GetRecommendationError](#) objects

instanceRecommendations

An array of objects that describe instance recommendations.

Type: Array of [InstanceRecommendation](#) objects

nextToken

The token to use to advance to the next page of instance recommendations.

This value is null when there are no more pages of instance recommendations to return.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

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

GetEC2RecommendationProjectedMetrics

Returns the projected utilization metrics of Amazon EC2 instance recommendations.

Note

The Cpu and Memory metrics are the only projected utilization metrics returned when you run this action. Additionally, the Memory metric is returned only for resources that have the unified CloudWatch agent installed on them. For more information, see [Enabling Memory Utilization with the CloudWatch Agent](#).

Request Syntax

```
{
  "endTime": number,
  "instanceArn": "string",
  "period": number,
  "recommendationPreferences": {
    "cpuVendorArchitectures": [ "string" ]
  },
  "startTime": number,
  "stat": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

endTime

The timestamp of the last projected metrics data point to return.

Type: Timestamp

Required: Yes

instanceArn

The Amazon Resource Name (ARN) of the instances for which to return recommendation projected metrics.

Type: String

Required: Yes

period

The granularity, in seconds, of the projected metrics data points.

Type: Integer

Required: Yes

recommendationPreferences

An object to specify the preferences for the Amazon EC2 recommendation projected metrics to return in the response.

Type: [RecommendationPreferences](#) object

Required: No

startTime

The timestamp of the first projected metrics data point to return.

Type: Timestamp

Required: Yes

stat

The statistic of the projected metrics.

Type: String

Valid Values: Maximum | Average

Required: Yes

Response Syntax

```
{
```

```
"recommendedOptionProjectedMetrics": [
  {
    "projectedMetrics": [
      {
        "name": "string",
        "timestamps": [ number ],
        "values": [ number ]
      }
    ],
    "rank": number,
    "recommendedInstanceType": "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.

recommendedOptionProjectedMetrics

An array of objects that describes projected metrics.

Type: Array of [RecommendedOptionProjectedMetric](#) objects

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerError

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetECSServiceRecommendationProjectedMetrics

Returns the projected metrics of Amazon ECS service recommendations.

Request Syntax

```
{  
  "endTime": number,  
  "period": number,  
  "serviceArn": "string",  
  "startTime": number,  
  "stat": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

endTime

The timestamp of the last projected metrics data point to return.

Type: Timestamp

Required: Yes

period

The granularity, in seconds, of the projected metrics data points.

Type: Integer

Required: Yes

serviceArn

The ARN that identifies the Amazon ECS service.

The following is the format of the ARN:

arn:aws:ecs:region:aws_account_id:service/cluster-name/service-name

Type: String

Required: Yes

startTime

The timestamp of the first projected metrics data point to return.

Type: Timestamp

Required: Yes

stat

The statistic of the projected metrics.

Type: String

Valid Values: Maximum | Average

Required: Yes

Response Syntax

```
{
  "recommendedOptionProjectedMetrics": [
    {
      "projectedMetrics": [
        {
          "lowerBoundValues": [ number ],
          "name": "string",
          "timestamps": [ number ],
          "upperBoundValues": [ number ]
        }
      ],
      "recommendedCpuUnits": number,
      "recommendedMemorySize": 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.

[recommendedOptionProjectedMetrics](#)

An array of objects that describes the projected metrics.

Type: Array of [ECSServiceRecommendedOptionProjectedMetric](#) objects

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetECSServiceRecommendations

Returns Amazon ECS service recommendations.

AWS Compute Optimizer generates recommendations for Amazon ECS services on Fargate that meet a specific set of requirements. For more information, see the [Supported resources and requirements](#) in the *AWS Compute Optimizer User Guide*.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "maxResults": number,
  "nextToken": "string",
  "serviceArns": [ "string" ]
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[accountIds](#)

Return the Amazon ECS service recommendations to the specified AWS account IDs.

If your account is the management account or the delegated administrator of an organization, use this parameter to return the Amazon ECS service recommendations to specific member accounts.

You can only specify one account ID per request.

Type: Array of strings

Required: No

filters

An array of objects to specify a filter that returns a more specific list of Amazon ECS service recommendations.

Type: Array of [ECSServiceRecommendationFilter](#) objects

Required: No

maxResults

The maximum number of Amazon ECS service recommendations to return with a single request.

To retrieve the remaining results, make another request with the returned `nextToken` value.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

nextToken

The token to advance to the next page of Amazon ECS service recommendations.

Type: String

Required: No

serviceArns

The ARN that identifies the Amazon ECS service.

The following is the format of the ARN:

`arn:aws:ecs:region:aws_account_id:service/cluster-name/service-name`

Type: Array of strings

Required: No

Response Syntax

```
{
  "ecsServiceRecommendations": [
    {
```

```

"accountId": "string",
"currentPerformanceRisk": "string",
"currentServiceConfiguration": {
  "autoScalingConfiguration": "string",
  "containerConfigurations": [
    {
      "containerName": "string",
      "cpu": number,
      "memorySizeConfiguration": {
        "memory": number,
        "memoryReservation": number
      }
    }
  ],
  "cpu": number,
  "memory": number,
  "taskDefinitionArn": "string"
},
"effectiveRecommendationPreferences": {
  "savingsEstimationMode": {
    "source": "string"
  }
},
"finding": "string",
"findingReasonCodes": [ "string" ],
"lastRefreshTimestamp": number,
"launchType": "string",
"lookbackPeriodInDays": number,
"serviceArn": "string",
"serviceRecommendationOptions": [
  {
    "containerRecommendations": [
      {
        "containerName": "string",
        "cpu": number,
        "memorySizeConfiguration": {
          "memory": number,
          "memoryReservation": number
        }
      }
    ]
  },
  "cpu": number,
  "memory": number,
  "projectedUtilizationMetrics": [

```

```
        {
            "lowerBoundValue": number,
            "name": "string",
            "statistic": "string",
            "upperBoundValue": number
        }
    ],
    "savingsOpportunity": {
        "estimatedMonthlySavings": {
            "currency": "string",
            "value": number
        },
        "savingsOpportunityPercentage": number
    },
    "savingsOpportunityAfterDiscounts": {
        "estimatedMonthlySavings": {
            "currency": "string",
            "value": number
        },
        "savingsOpportunityPercentage": number
    }
}
],
"tags": [
    {
        "key": "string",
        "value": "string"
    }
],
"utilizationMetrics": [
    {
        "name": "string",
        "statistic": "string",
        "value": number
    }
]
}
],
"errors": [
    {
        "code": "string",
        "identifier": "string",
        "message": "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.

ecsServiceRecommendations

An array of objects that describe the Amazon ECS service recommendations.

Type: Array of [ECSServiceRecommendation](#) objects

errors

An array of objects that describe errors of the request.

Type: Array of [GetRecommendationError](#) objects

nextToken

The token to advance to the next page of Amazon ECS service recommendations.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetEffectiveRecommendationPreferences

Returns the recommendation preferences that are in effect for a given resource, such as enhanced infrastructure metrics. Considers all applicable preferences that you might have set at the resource, account, and organization level.

When you create a recommendation preference, you can set its status to Active or Inactive. Use this action to view the recommendation preferences that are in effect, or Active.

Request Syntax

```
{  
  "resourceArn": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

resourceArn

The Amazon Resource Name (ARN) of the resource for which to confirm effective recommendation preferences. Only EC2 instance and Auto Scaling group ARNs are currently supported.

Type: String

Required: Yes

Response Syntax

```
{  
  "enhancedInfrastructureMetrics": "string",  
  "externalMetricsPreference": {  
    "source": "string"  
  },  
  "lookBackPeriod": "string",  
  "preferredResources": [  
    {  
      "resourceArn": "string",  
      "status": "string"  
    }  
  ]  
}
```



```
{
  "effectiveIncludeList": [ "string" ],
  "excludeList": [ "string" ],
  "includeList": [ "string" ],
  "name": "string"
},
"utilizationPreferences": [
  {
    "metricName": "string",
    "metricParameters": {
      "headroom": "string",
      "threshold": "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.

enhancedInfrastructureMetrics

The status of the enhanced infrastructure metrics recommendation preference. Considers all applicable preferences that you might have set at the resource, account, and organization level.

A status of `Active` confirms that the preference is applied in the latest recommendation refresh, and a status of `Inactive` confirms that it's not yet applied to recommendations.

To validate whether the preference is applied to your last generated set of recommendations, review the `effectiveRecommendationPreferences` value in the response of the [GetAutoScalingGroupRecommendations](#) and [GetEC2InstanceRecommendations](#) actions.

For more information, see [Enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Type: String

Valid Values: `Active` | `Inactive`

externalMetricsPreference

The provider of the external metrics recommendation preference. Considers all applicable preferences that you might have set at the account and organization level.

If the preference is applied in the latest recommendation refresh, an object with a valid source value appears in the response. If the preference isn't applied to the recommendations already, then this object doesn't appear in the response.

To validate whether the preference is applied to your last generated set of recommendations, review the `effectiveRecommendationPreferences` value in the response of the [GetEC2InstanceRecommendations](#) actions.

For more information, see [Enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Type: [ExternalMetricsPreference](#) object

lookBackPeriod

The number of days the utilization metrics of the AWS resource are analyzed.

To validate that the preference is applied to your last generated set of recommendations, review the `effectiveRecommendationPreferences` value in the response of the `GetAutoScalingGroupRecommendations` or `GetEC2InstanceRecommendations` actions.

Type: String

Valid Values: DAYS_14 | DAYS_32 | DAYS_93

preferredResources

The resource type values that are considered as candidates when generating rightsizing recommendations. This object resolves any wildcard expressions and returns the effective list of candidate resource type values. It also considers all applicable preferences that you set at the resource, account, and organization level.

To validate that the preference is applied to your last generated set of recommendations, review the `effectiveRecommendationPreferences` value in the response of the `GetAutoScalingGroupRecommendations` or `GetEC2InstanceRecommendations` actions.

Type: Array of [EffectivePreferredResource](#) objects

utilizationPreferences

The resource's CPU and memory utilization preferences, such as threshold and headroom, that were used to generate rightsizing recommendations. It considers all applicable preferences that you set at the resource, account, and organization level.

To validate that the preference is applied to your last generated set of recommendations, review the `effectiveRecommendationPreferences` value in the response of the `GetAutoScalingGroupRecommendations` or `GetEC2InstanceRecommendations` actions.

Type: Array of [UtilizationPreference](#) objects

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetEnrollmentStatus

Returns the enrollment (opt in) status of an account to the AWS Compute Optimizer service.

If the account is the management account of an organization, this action also confirms the enrollment status of member accounts of the organization. Use the [GetEnrollmentStatusesForOrganization](#) action to get detailed information about the enrollment status of member accounts of an organization.

Response Syntax

```
{
  "lastUpdatedTimestamp": number,
  "memberAccountsEnrolled": boolean,
  "numberOfMemberAccountsOptedIn": number,
  "status": "string",
  "statusReason": "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.

[lastUpdatedTimestamp](#)

The Unix epoch timestamp, in seconds, of when the account enrollment status was last updated.

Type: Timestamp

[memberAccountsEnrolled](#)

Confirms the enrollment status of member accounts of the organization, if the account is a management account of an organization.

Type: Boolean

[numberOfMemberAccountsOptedIn](#)

The count of organization member accounts that are opted in to the service, if your account is an organization management account.

Type: Integer

status

The enrollment status of the account.

Type: String

Valid Values: Active | Inactive | Pending | Failed

statusReason

The reason for the enrollment status of the account.

For example, an account might show a status of Pending because member accounts of an organization require more time to be enrolled in the service.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetEnrollmentStatusesForOrganization

Returns the AWS Compute Optimizer enrollment (opt-in) status of organization member accounts, if your account is an organization management account.

To get the enrollment status of standalone accounts, use the [GetEnrollmentStatus](#) action.

Request Syntax

```
{
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "maxResults": number,
  "nextToken": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[filters](#)

An array of objects to specify a filter that returns a more specific list of account enrollment statuses.

Type: Array of [EnrollmentFilter](#) objects

Required: No

[maxResults](#)

The maximum number of account enrollment statuses to return with a single request. You can specify up to 100 statuses to return with each request.

To retrieve the remaining results, make another request with the returned nextToken value.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

[nextToken](#)

The token to advance to the next page of account enrollment statuses.

Type: String

Required: No

Response Syntax

```
{
  "accountEnrollmentStatuses": [
    {
      "accountId": "string",
      "lastUpdatedTimestamp": number,
      "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.

[accountEnrollmentStatuses](#)

An array of objects that describe the enrollment statuses of organization member accounts.

Type: Array of [AccountEnrollmentStatus](#) objects

[nextToken](#)

The token to use to advance to the next page of account enrollment statuses.

This value is null when there are no more pages of account enrollment statuses to return.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetIdleRecommendations

Returns idle resource recommendations. AWS Compute Optimizer generates recommendations for idle resources that meet a specific set of requirements. For more information, see [Resource requirements](#) in the *AWS Compute Optimizer User Guide*

Request Syntax

```
{
  "accountIds": [ "string" ],
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "maxResults": number,
  "nextToken": "string",
  "orderBy": {
    "dimension": "string",
    "order": "string"
  },
  "resourceArns": [ "string" ]
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[accountIds](#)

Return the idle resource recommendations to the specified AWS account IDs.

If your account is the management account or the delegated administrator of an organization, use this parameter to return the idle resource recommendations to specific member accounts.

You can only specify one account ID per request.

Type: Array of strings

Required: No

filters

An array of objects to specify a filter that returns a more specific list of idle resource recommendations.

Type: Array of [IdleRecommendationFilter](#) objects

Required: No

maxResults

The maximum number of idle resource recommendations to return with a single request.

To retrieve the remaining results, make another request with the returned `nextToken` value.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 100.

Required: No

nextToken

The token to advance to the next page of idle resource recommendations.

Type: String

Required: No

orderBy

The order to sort the idle resource recommendations.

Type: [OrderBy](#) object

Required: No

resourceArns

The ARN that identifies the idle resource.

Type: Array of strings

Required: No

Response Syntax

```
{
  "errors": [
    {
      "code": "string",
      "identifier": "string",
      "message": "string",
      "resourceType": "string"
    }
  ],
  "idleRecommendations": [
    {
      "accountId": "string",
      "finding": "string",
      "findingDescription": "string",
      "lastRefreshTimestamp": number,
      "lookBackPeriodInDays": number,
      "resourceArn": "string",
      "resourceId": "string",
      "resourceType": "string",
      "savingsOpportunity": {
        "estimatedMonthlySavings": {
          "currency": "string",
          "value": number
        },
        "savingsOpportunityPercentage": number
      },
      "savingsOpportunityAfterDiscounts": {
        "estimatedMonthlySavings": {
          "currency": "string",
          "value": number
        },
        "savingsOpportunityPercentage": number
      },
      "tags": [
        {
          "key": "string",
          "value": "string"
        }
      ],
      "utilizationMetrics": [
        {
```

```
        "name": "string",
        "statistic": "string",
        "value": 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.

errors

An array of objects that describe errors of the request.

Type: Array of [IdleRecommendationError](#) objects

idleRecommendations

An array of objects that describe the idle resource recommendations.

Type: Array of [IdleRecommendation](#) objects

nextToken

The token to advance to the next page of idle resource recommendations.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetLambdaFunctionRecommendations

Returns AWS Lambda function recommendations.

AWS Compute Optimizer generates recommendations for functions that meet a specific set of requirements. For more information, see the [Supported resources and requirements](#) in the *AWS Compute Optimizer User Guide*.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "functionArns": [ "string" ],
  "maxResults": number,
  "nextToken": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[accountIds](#)

The ID of the AWS account for which to return function recommendations.

If your account is the management account of an organization, use this parameter to specify the member account for which you want to return function recommendations.

Only one account ID can be specified per request.

Type: Array of strings

Required: No

filters

An array of objects to specify a filter that returns a more specific list of function recommendations.

Type: Array of [LambdaFunctionRecommendationFilter](#) objects

Required: No

functionArns

The Amazon Resource Name (ARN) of the functions for which to return recommendations.

You can specify a qualified or unqualified ARN. If you specify an unqualified ARN without a function version suffix, Compute Optimizer will return recommendations for the latest (\$LATEST) version of the function. If you specify a qualified ARN with a version suffix, Compute Optimizer will return recommendations for the specified function version. For more information about using function versions, see [Using versions](#) in the *AWS Lambda Developer Guide*.

Type: Array of strings

Required: No

maxResults

The maximum number of function recommendations to return with a single request.

To retrieve the remaining results, make another request with the returned nextToken value.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

nextToken

The token to advance to the next page of function recommendations.

Type: String

Required: No

Response Syntax

```
{
  "lambdaFunctionRecommendations": [
    {
      "accountId": "string",
      "currentMemorySize": number,
      "currentPerformanceRisk": "string",
      "effectiveRecommendationPreferences": {
        "savingsEstimationMode": {
          "source": "string"
        }
      },
      "finding": "string",
      "findingReasonCodes": [ "string" ],
      "functionArn": "string",
      "functionVersion": "string",
      "lastRefreshTimestamp": number,
      "lookbackPeriodInDays": number,
      "memorySizeRecommendationOptions": [
        {
          "memorySize": number,
          "projectedUtilizationMetrics": [
            {
              "name": "string",
              "statistic": "string",
              "value": number
            }
          ],
          "rank": number,
          "savingsOpportunity": {
            "estimatedMonthlySavings": {
              "currency": "string",
              "value": number
            },
            "savingsOpportunityPercentage": number
          },
          "savingsOpportunityAfterDiscounts": {
            "estimatedMonthlySavings": {
              "currency": "string",
              "value": number
            },
            "savingsOpportunityPercentage": number
          }
        }
      ]
    }
  ]
}
```

```
    }
  }
],
"numberOfInvocations": number,
"tags": [
  {
    "key": "string",
    "value": "string"
  }
],
"utilizationMetrics": [
  {
    "name": "string",
    "statistic": "string",
    "value": 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.

lambdaFunctionRecommendations

An array of objects that describe function recommendations.

Type: Array of [LambdaFunctionRecommendation](#) objects

nextToken

The token to use to advance to the next page of function recommendations.

This value is null when there are no more pages of function recommendations to return.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

LimitExceededException

The request exceeds a limit of the service.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetLicenseRecommendations

Returns license recommendations for Amazon EC2 instances that run on a specific license.

AWS Compute Optimizer generates recommendations for licenses that meet a specific set of requirements. For more information, see the [Supported resources and requirements](#) in the *AWS Compute Optimizer User Guide*.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "maxResults": number,
  "nextToken": "string",
  "resourceArns": [ "string" ]
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

accountIds

The ID of the AWS account for which to return license recommendations.

If your account is the management account of an organization, use this parameter to specify the member account for which you want to return license recommendations.

Only one account ID can be specified per request.

Type: Array of strings

Required: No

filters

An array of objects to specify a filter that returns a more specific list of license recommendations.

Type: Array of [LicenseRecommendationFilter](#) objects

Required: No

maxResults

The maximum number of license recommendations to return with a single request.

To retrieve the remaining results, make another request with the returned `nextToken` value.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

nextToken

The token to advance to the next page of license recommendations.

Type: String

Required: No

resourceArns

The ARN that identifies the Amazon EC2 instance.

The following is the format of the ARN:

```
arn:aws:ec2:region:aws_account_id:instance/instance-id
```

Type: Array of strings

Required: No

Response Syntax

```
{  
  "errors": [  
    {
```

```

        "code": "string",
        "identifier": "string",
        "message": "string"
    }
],
"licenseRecommendations": [
    {
        "accountId": "string",
        "currentLicenseConfiguration": {
            "instanceType": "string",
            "licenseEdition": "string",
            "licenseModel": "string",
            "licenseName": "string",
            "licenseVersion": "string",
            "metricsSource": [
                {
                    "provider": "string",
                    "providerArn": "string"
                }
            ],
            "numberOfCores": number,
            "operatingSystem": "string"
        },
        "finding": "string",
        "findingReasonCodes": [ "string" ],
        "lastRefreshTimestamp": number,
        "licenseRecommendationOptions": [
            {
                "licenseEdition": "string",
                "licenseModel": "string",
                "operatingSystem": "string",
                "rank": number,
                "savingsOpportunity": {
                    "estimatedMonthlySavings": {
                        "currency": "string",
                        "value": number
                    },
                    "savingsOpportunityPercentage": number
                }
            }
        ],
        "lookbackPeriodInDays": number,
        "resourceArn": "string",
        "tags": [

```

```
{
  {
    "key": "string",
    "value": "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.

errors

An array of objects that describe errors of the request.

Type: Array of [GetRecommendationError](#) objects

licenseRecommendations

An array of objects that describe license recommendations.

Type: Array of [LicenseRecommendation](#) objects

nextToken

The token to use to advance to the next page of license recommendations.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetRDSDatabaseRecommendationProjectedMetrics

Returns the projected metrics of Aurora and RDS database recommendations.

Request Syntax

```
{
  "endTime": number,
  "period": number,
  "recommendationPreferences": {
    "cpuVendorArchitectures": [ "string" ]
  },
  "resourceArn": "string",
  "startTime": number,
  "stat": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

endTime

The timestamp of the last projected metrics data point to return.

Type: Timestamp

Required: Yes

period

The granularity, in seconds, of the projected metrics data points.

Type: Integer

Required: Yes

recommendationPreferences

Describes the recommendation preferences to return in the response of a [GetAutoScalingGroupRecommendations](#), [GetEC2InstanceRecommendations](#),

[GetEC2RecommendationProjectedMetrics](#), [GetRDSDatabaseRecommendations](#), and [GetRDSDatabaseRecommendationProjectedMetrics](#) request.

Type: [RecommendationPreferences](#) object

Required: No

[resourceArn](#)

The ARN that identifies the Amazon Aurora or RDS database.

The following is the format of the ARN:

```
arn:aws:rds:{region}:{accountId}:db:{resourceName}
```

Type: String

Required: Yes

[startTime](#)

The timestamp of the first projected metrics data point to return.

Type: Timestamp

Required: Yes

[stat](#)

The statistic of the projected metrics.

Type: String

Valid Values: Maximum | Average

Required: Yes

Response Syntax

```
{
  "recommendedOptionProjectedMetrics": [
    {
      "projectedMetrics": [
        {
          "name": "string",
```

```
        "timestamps": [ number ],
        "values": [ number ]
      }
    ],
    "rank": number,
    "recommendedDBInstanceClass": "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.

recommendedOptionProjectedMetrics

An array of objects that describes the projected metrics.

Type: Array of [RDSDatabaseRecommendedOptionProjectedMetric](#) objects

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerError

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)

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

GetRDSDatabaseRecommendations

Returns Amazon Aurora and RDS database recommendations.

AWS Compute Optimizer generates recommendations for Amazon Aurora and RDS databases that meet a specific set of requirements. For more information, see the [Supported resources and requirements](#) in the *AWS Compute Optimizer User Guide*.

Request Syntax

```
{
  "accountIds": [ "string" ],
  "filters": [
    {
      "name": "string",
      "values": [ "string" ]
    }
  ],
  "maxResults": number,
  "nextToken": "string",
  "recommendationPreferences": {
    "cpuVendorArchitectures": [ "string" ]
  },
  "resourceArns": [ "string" ]
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[accountIds](#)

Return the Amazon Aurora and RDS database recommendations to the specified AWS account IDs.

If your account is the management account or the delegated administrator of an organization, use this parameter to return the Amazon Aurora and RDS database recommendations to specific member accounts.

You can only specify one account ID per request.

Type: Array of strings

Required: No

filters

An array of objects to specify a filter that returns a more specific list of Amazon Aurora and RDS database recommendations.

Type: Array of [RDSDBRecommendationFilter](#) objects

Required: No

maxResults

The maximum number of Amazon Aurora and RDS database recommendations to return with a single request.

To retrieve the remaining results, make another request with the returned `nextToken` value.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

nextToken

The token to advance to the next page of Amazon Aurora and RDS database recommendations.

Type: String

Required: No

recommendationPreferences

Describes the recommendation preferences to return in the response of a [GetAutoScalingGroupRecommendations](#), [GetEC2InstanceRecommendations](#), [GetEC2RecommendationProjectedMetrics](#), [GetRDSDatabaseRecommendations](#), and [GetRDSDatabaseRecommendationProjectedMetrics](#) request.

Type: [RecommendationPreferences](#) object

Required: No

resourceArns

The ARN that identifies the Amazon Aurora or RDS database.

The following is the format of the ARN:

```
arn:aws:rds:{region}:{accountId}:db:{resourceName}
```

The following is the format of a DB Cluster ARN:

```
arn:aws:rds:{region}:{accountId}:cluster:{resourceName}
```

Type: Array of strings

Required: No

Response Syntax

```
{
  "errors": [
    {
      "code": "string",
      "identifier": "string",
      "message": "string"
    }
  ],
  "nextToken": "string",
  "rdsDBRecommendations": [
    {
      "accountId": "string",
      "currentDBInstanceClass": "string",
      "currentInstancePerformanceRisk": "string",
      "currentStorageConfiguration": {
        "allocatedStorage": number,
        "iops": number,
        "maxAllocatedStorage": number,
        "storageThroughput": number,
        "storageType": "string"
      },
      "currentStorageEstimatedMonthlyVolumeIOPsCostVariation": "string",
      "dbClusterIdentifier": "string",
      "effectiveRecommendationPreferences": {
        "cpuVendorArchitectures": [ "string" ],
        "enhancedInfrastructureMetrics": "string",
        "lookBackPeriod": "string",
        "savingsEstimationMode": {
          "source": "string"
        }
      }
    }
  ]
}
```

```

    }
  },
  "engine": "string",
  "engineVersion": "string",
  "idle": "string",
  "instanceFinding": "string",
  "instanceFindingReasonCodes": [ "string" ],
  "instanceRecommendationOptions": [
    {
      "dbInstanceClass": "string",
      "performanceRisk": number,
      "projectedUtilizationMetrics": [
        {
          "name": "string",
          "statistic": "string",
          "value": number
        }
      ],
      "rank": number,
      "savingsOpportunity": {
        "estimatedMonthlySavings": {
          "currency": "string",
          "value": number
        },
        "savingsOpportunityPercentage": number
      },
      "savingsOpportunityAfterDiscounts": {
        "estimatedMonthlySavings": {
          "currency": "string",
          "value": number
        },
        "savingsOpportunityPercentage": number
      }
    }
  ],
  "lastRefreshTimestamp": number,
  "lookbackPeriodInDays": number,
  "promotionTier": number,
  "resourceArn": "string",
  "storageFinding": "string",
  "storageFindingReasonCodes": [ "string" ],
  "storageRecommendationOptions": [
    {
      "estimatedMonthlyVolumeIOPsCostVariation": "string",

```

```

    "rank": number,
    "savingsOpportunity": {
      "estimatedMonthlySavings": {
        "currency": "string",
        "value": number
      },
      "savingsOpportunityPercentage": number
    },
    "savingsOpportunityAfterDiscounts": {
      "estimatedMonthlySavings": {
        "currency": "string",
        "value": number
      },
      "savingsOpportunityPercentage": number
    },
    "storageConfiguration": {
      "allocatedStorage": number,
      "iops": number,
      "maxAllocatedStorage": number,
      "storageThroughput": number,
      "storageType": "string"
    }
  ],
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ],
  "utilizationMetrics": [
    {
      "name": "string",
      "statistic": "string",
      "value": 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.

errors

An array of objects that describe errors of the request.

Type: Array of [GetRecommendationError](#) objects

nextToken

The token to advance to the next page of Amazon Aurora and RDS database recommendations.

Type: String

rdsDBRecommendations

An array of objects that describe the Amazon Aurora and RDS database recommendations.

Type: Array of [RDSDBRecommendation](#) objects

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerError

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)

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

GetRecommendationPreferences

Returns existing recommendation preferences, such as enhanced infrastructure metrics.

Use the `scope` parameter to specify which preferences to return. You can specify to return preferences for an organization, a specific account ID, or a specific EC2 instance or Auto Scaling group Amazon Resource Name (ARN).

For more information, see [Activating enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Request Syntax

```
{
  "maxResults": number,
  "nextToken": "string",
  "resourceType": "string",
  "scope": {
    "name": "string",
    "value": "string"
  }
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[maxResults](#)

The maximum number of recommendation preferences to return with a single request.

To retrieve the remaining results, make another request with the returned `nextToken` value.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

nextToken

The token to advance to the next page of recommendation preferences.

Type: String

Required: No

resourceType

The target resource type of the recommendation preference for which to return preferences.

The `Ec2Instance` option encompasses standalone instances and instances that are part of Auto Scaling groups. The `AutoScalingGroup` option encompasses only instances that are part of an Auto Scaling group.

Type: String

Valid Values: `Ec2Instance` | `AutoScalingGroup` | `RdsDBInstance` | `AuroraDBClusterStorage`

Required: Yes

scope

An object that describes the scope of the recommendation preference to return.

You can return recommendation preferences that are created at the organization level (for management accounts of an organization only), account level, and resource level. For more information, see [Activating enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Type: [Scope](#) object

Required: No

Response Syntax

```
{
  "nextToken": "string",
  "recommendationPreferencesDetails": [
    {
      "enhancedInfrastructureMetrics": "string",
```

```

    "externalMetricsPreference": {
      "source": "string"
    },
    "inferredWorkloadTypes": "string",
    "lookBackPeriod": "string",
    "preferredResources": [
      {
        "effectiveIncludeList": [ "string" ],
        "excludeList": [ "string" ],
        "includeList": [ "string" ],
        "name": "string"
      }
    ],
    "resourceType": "string",
    "savingsEstimationMode": "string",
    "scope": {
      "name": "string",
      "value": "string"
    },
    "utilizationPreferences": [
      {
        "metricName": "string",
        "metricParameters": {
          "headroom": "string",
          "threshold": "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.

[nextToken](#)

The token to use to advance to the next page of recommendation preferences.

This value is null when there are no more pages of recommendation preferences to return.

Type: String

recommendationPreferencesDetails

An array of objects that describe recommendation preferences.

Type: Array of [RecommendationPreferencesDetail](#) objects

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerError

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

GetRecommendationSummaries

Returns the optimization findings for an account.

It returns the number of:

- Amazon EC2 instances in an account that are Underprovisioned, Overprovisioned, or Optimized.
- EC2Auto Scaling groups in an account that are NotOptimized, or Optimized.
- Amazon EBS volumes in an account that are NotOptimized, or Optimized.
- Lambda functions in an account that are NotOptimized, or Optimized.
- Amazon ECS services in an account that are Underprovisioned, Overprovisioned, or Optimized.
- Commercial software licenses in an account that are InsufficientMetrics, NotOptimized or Optimized.
- Amazon Aurora and Amazon RDS databases in an account that are Underprovisioned, Overprovisioned, Optimized, or NotOptimized.

Request Syntax

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

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

accountIds

The ID of the AWS account for which to return recommendation summaries.

If your account is the management account of an organization, use this parameter to specify the member account for which you want to return recommendation summaries.

Only one account ID can be specified per request.

Type: Array of strings

Required: No

maxResults

The maximum number of recommendation summaries to return with a single request.

To retrieve the remaining results, make another request with the returned `nextToken` value.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

nextToken

The token to advance to the next page of recommendation summaries.

Type: String

Required: No

Response Syntax

```
{
  "nextToken": "string",
  "recommendationSummaries": [
    {
      "accountId": "string",
      "aggregatedSavingsOpportunity": {
        "estimatedMonthlySavings": {
          "currency": "string",
          "value": number
        },
        "savingsOpportunityPercentage": number
      },
      "currentPerformanceRiskRatings": {
        "high": number,
        "low": number,
        "medium": number,
        "veryLow": number
      }
    }
  ]
}
```

```

    },
    "idleSavingsOpportunity": {
      "estimatedMonthlySavings": {
        "currency": "string",
        "value": number
      },
      "savingsOpportunityPercentage": number
    },
    "idleSummaries": [
      {
        "name": "string",
        "value": number
      }
    ],
    "inferredWorkloadSavings": [
      {
        "estimatedMonthlySavings": {
          "currency": "string",
          "value": number
        },
        "inferredWorkloadTypes": [ "string" ]
      }
    ],
    "recommendationResourceType": "string",
    "savingsOpportunity": {
      "estimatedMonthlySavings": {
        "currency": "string",
        "value": number
      },
      "savingsOpportunityPercentage": number
    },
    "summaries": [
      {
        "name": "string",
        "reasonCodeSummaries": [
          {
            "name": "string",
            "value": number
          }
        ],
        "value": 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.

[nextToken](#)

The token to use to advance to the next page of recommendation summaries.

This value is null when there are no more pages of recommendation summaries to return.

Type: String

[recommendationSummaries](#)

An array of objects that summarize a recommendation.

Type: Array of [RecommendationSummary](#) objects

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerError

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

PutRecommendationPreferences

Creates a new recommendation preference or updates an existing recommendation preference, such as enhanced infrastructure metrics.

For more information, see [Activating enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Request Syntax

```
{
  "enhancedInfrastructureMetrics": "string",
  "externalMetricsPreference": {
    "source": "string"
  },
  "inferredWorkloadTypes": "string",
  "lookBackPeriod": "string",
  "preferredResources": [
    {
      "excludeList": [ "string" ],
      "includeList": [ "string" ],
      "name": "string"
    }
  ],
  "resourceType": "string",
  "savingsEstimationMode": "string",
  "scope": {
    "name": "string",
    "value": "string"
  },
  "utilizationPreferences": [
    {
      "metricName": "string",
      "metricParameters": {
        "headroom": "string",
        "threshold": "string"
      }
    }
  ]
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[enhancedInfrastructureMetrics](#)

The status of the enhanced infrastructure metrics recommendation preference to create or update.

Specify the `Active` status to activate the preference, or specify `Inactive` to deactivate the preference.

For more information, see [Enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Type: String

Valid Values: `Active` | `Inactive`

Required: No

[externalMetricsPreference](#)

The provider of the external metrics recommendation preference to create or update.

Specify a valid provider in the `source` field to activate the preference. To delete this preference, see the [DeleteRecommendationPreferences](#) action.

This preference can only be set for the `Ec2Instance` resource type.

For more information, see [External metrics ingestion](#) in the *AWS Compute Optimizer User Guide*.

Type: [ExternalMetricsPreference](#) object

Required: No

[inferredWorkloadTypes](#)

The status of the inferred workload types recommendation preference to create or update.

Note

The inferred workload type feature is active by default. To deactivate it, create a recommendation preference.

Specify the `Inactive` status to deactivate the feature, or specify `Active` to activate it.

For more information, see [Inferred workload types](#) in the *AWS Compute Optimizer User Guide*.

Type: String

Valid Values: `Active` | `Inactive`

Required: No

lookBackPeriod

The preference to control the number of days the utilization metrics of the AWS resource are analyzed. When this preference isn't specified, we use the default value `DAYS_14`.

You can only set this preference for the Amazon EC2 instance and Auto Scaling group resource types.

Note

- Amazon EC2 instance lookback preferences can be set at the organization, account, and resource levels.
- Auto Scaling group lookback preferences can only be set at the resource level.

Type: String

Valid Values: `DAYS_14` | `DAYS_32` | `DAYS_93`

Required: No

preferredResources

The preference to control which resource type values are considered when generating rightsizing recommendations. You can specify this preference as a combination of include and

exclude lists. You must specify either an `includeList` or `excludeList`. If the preference is an empty set of resource type values, an error occurs.

 Note

You can only set this preference for the Amazon EC2 instance and Auto Scaling group resource types.

Type: Array of [PreferredResource](#) objects

Required: No

[resourceType](#)

The target resource type of the recommendation preference to create.

The `Ec2Instance` option encompasses standalone instances and instances that are part of Auto Scaling groups. The `AutoScalingGroup` option encompasses only instances that are part of an Auto Scaling group.

Type: String

Valid Values: `Ec2Instance` | `AutoScalingGroup` | `RdsDBInstance` | `AuroraDBClusterStorage`

Required: Yes

[savingsEstimationMode](#)

The status of the savings estimation mode preference to create or update.

Specify the `AfterDiscounts` status to activate the preference, or specify `BeforeDiscounts` to deactivate the preference.

Only the account manager or delegated administrator of your organization can activate this preference.

For more information, see [Savings estimation mode](#) in the *AWS Compute Optimizer User Guide*.

Type: String

Valid Values: AfterDiscounts | BeforeDiscounts

Required: No

scope

An object that describes the scope of the recommendation preference to create.

You can create recommendation preferences at the organization level (for management accounts of an organization only), account level, and resource level. For more information, see [Activating enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Note

You cannot create recommendation preferences for Auto Scaling groups at the organization and account levels. You can create recommendation preferences for Auto Scaling groups only at the resource level by specifying a scope name of ResourceArn and a scope value of the Auto Scaling group Amazon Resource Name (ARN). This will configure the preference for all instances that are part of the specified Auto Scaling group. You also cannot create recommendation preferences at the resource level for instances that are part of an Auto Scaling group. You can create recommendation preferences at the resource level only for standalone instances.

Type: [Scope](#) object

Required: No

utilizationPreferences

The preference to control the resource's CPU utilization threshold, CPU utilization headroom, and memory utilization headroom. When this preference isn't specified, we use the following default values.

CPU utilization:

- P99_5 for threshold
- PERCENT_20 for headroom

Memory utilization:

- PERCENT_20 for headroom

Note

- You can only set CPU and memory utilization preferences for the Amazon EC2 instance resource type.
- The threshold setting isn't available for memory utilization.

Type: Array of [UtilizationPreference](#) objects

Required: No

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

OptInRequiredException

The account is not opted in to AWS Compute Optimizer.

HTTP Status Code: 400

ResourceNotFoundException

A resource that is required for the action doesn't exist.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

UpdateEnrollmentStatus

Updates the enrollment (opt in and opt out) status of an account to the AWS Compute Optimizer service.

If the account is a management account of an organization, this action can also be used to enroll member accounts of the organization.

You must have the appropriate permissions to opt in to Compute Optimizer, to view its recommendations, and to opt out. For more information, see [Controlling access with AWS Identity and Access Management](#) in the *AWS Compute Optimizer User Guide*.

When you opt in, Compute Optimizer automatically creates a service-linked role in your account to access its data. For more information, see [Using Service-Linked Roles for AWS Compute Optimizer](#) in the *AWS Compute Optimizer User Guide*.

Request Syntax

```
{
  "includeMemberAccounts": boolean,
  "status": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#).

The request accepts the following data in JSON format.

[includeMemberAccounts](#)

Indicates whether to enroll member accounts of the organization if the account is the management account of an organization.

Type: Boolean

Required: No

[status](#)

The new enrollment status of the account.

The following status options are available:

- **Active** - Opt in your account to the Compute Optimizer service. Compute Optimizer begins analyzing the configuration and utilization metrics of your AWS resources after you opt in. For more information, see [Metrics analyzed by AWS Compute Optimizer](#) in the *AWS Compute Optimizer User Guide*.
- **Inactive** - Opt out your account from the Compute Optimizer service. Your account's recommendations and related metrics data will be deleted from Compute Optimizer after you opt out.

 Note

The Pending and Failed options cannot be used to update the enrollment status of an account. They are returned in the response of a request to update the enrollment status of an account.

Type: String

Valid Values: Active | Inactive | Pending | Failed

Required: Yes

Response Syntax

```
{
  "status": "string",
  "statusReason": "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.

status

The enrollment status of the account.

Type: String

Valid Values: Active | Inactive | Pending | Failed

statusReason

The reason for the enrollment status of the account. For example, an account might show a status of Pending because member accounts of an organization require more time to be enrolled in the service.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors](#).

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

InternalServerErrorException

An internal error has occurred. Try your call again.

HTTP Status Code: 500

InvalidParameterValueException

The value supplied for the input parameter is out of range or not valid.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 400

ServiceUnavailableException

The request has failed due to a temporary failure of the server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

Data Types

The AWS Compute Optimizer 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:

- [AccountEnrollmentStatus](#)
- [AutoScalingGroupConfiguration](#)
- [AutoScalingGroupEstimatedMonthlySavings](#)
- [AutoScalingGroupRecommendation](#)
- [AutoScalingGroupRecommendationOption](#)
- [AutoScalingGroupSavingsOpportunityAfterDiscounts](#)
- [ContainerConfiguration](#)
- [ContainerRecommendation](#)
- [CurrentPerformanceRiskRatings](#)
- [CustomizableMetricParameters](#)
- [DBStorageConfiguration](#)
- [EBSEffectiveRecommendationPreferences](#)
- [EBSEstimatedMonthlySavings](#)
- [EBSFilter](#)
- [EBSSavingsEstimationMode](#)
- [EBSSavingsOpportunityAfterDiscounts](#)
- [EBSUtilizationMetric](#)
- [ECSEffectiveRecommendationPreferences](#)
- [ECSEstimatedMonthlySavings](#)
- [ECSSavingsEstimationMode](#)

- [ECSSavingsOpportunityAfterDiscounts](#)
- [ECSServiceProjectedMetric](#)
- [ECSServiceProjectedUtilizationMetric](#)
- [ECSServiceRecommendation](#)
- [ECSServiceRecommendationFilter](#)
- [ECSServiceRecommendationOption](#)
- [ECSServiceRecommendedOptionProjectedMetric](#)
- [ECSServiceUtilizationMetric](#)
- [EffectivePreferredResource](#)
- [EffectiveRecommendationPreferences](#)
- [EnrollmentFilter](#)
- [EstimatedMonthlySavings](#)
- [ExportDestination](#)
- [ExternalMetricsPreference](#)
- [ExternalMetricStatus](#)
- [Filter](#)
- [GetRecommendationError](#)
- [Gpu](#)
- [GpuInfo](#)
- [IdleEstimatedMonthlySavings](#)
- [IdleRecommendation](#)
- [IdleRecommendationError](#)
- [IdleRecommendationFilter](#)
- [IdleSavingsOpportunity](#)
- [IdleSavingsOpportunityAfterDiscounts](#)
- [IdleSummary](#)
- [IdleUtilizationMetric](#)
- [InferredWorkloadSaving](#)
- [InstanceEstimatedMonthlySavings](#)
- [InstanceRecommendation](#)

- [InstanceRecommendationOption](#)
- [InstanceSavingsEstimationMode](#)
- [InstanceSavingsOpportunityAfterDiscounts](#)
- [JobFilter](#)
- [LambdaEffectiveRecommendationPreferences](#)
- [LambdaEstimatedMonthlySavings](#)
- [LambdaFunctionMemoryProjectedMetric](#)
- [LambdaFunctionMemoryRecommendationOption](#)
- [LambdaFunctionRecommendation](#)
- [LambdaFunctionRecommendationFilter](#)
- [LambdaFunctionUtilizationMetric](#)
- [LambdaSavingsEstimationMode](#)
- [LambdaSavingsOpportunityAfterDiscounts](#)
- [LicenseConfiguration](#)
- [LicenseRecommendation](#)
- [LicenseRecommendationFilter](#)
- [LicenseRecommendationOption](#)
- [MemorySizeConfiguration](#)
- [MetricSource](#)
- [OrderBy](#)
- [PreferredResource](#)
- [ProjectedMetric](#)
- [RDSDatabaseProjectedMetric](#)
- [RDSDatabaseRecommendedOptionProjectedMetric](#)
- [RDSDBInstanceRecommendationOption](#)
- [RDSDBRecommendation](#)
- [RDSDBRecommendationFilter](#)
- [RDSDBStorageRecommendationOption](#)
- [RDSDBUtilizationMetric](#)
- [RDSEffectiveRecommendationPreferences](#)

- [RDSInstanceEstimatedMonthlySavings](#)
- [RDSInstanceSavingsOpportunityAfterDiscounts](#)
- [RDSSavingsEstimationMode](#)
- [RDSStorageEstimatedMonthlySavings](#)
- [RDSStorageSavingsOpportunityAfterDiscounts](#)
- [ReasonCodeSummary](#)
- [RecommendationExportJob](#)
- [RecommendationPreferences](#)
- [RecommendationPreferencesDetail](#)
- [RecommendationSource](#)
- [RecommendationSummary](#)
- [RecommendedOptionProjectedMetric](#)
- [S3Destination](#)
- [S3DestinationConfig](#)
- [SavingsOpportunity](#)
- [Scope](#)
- [ServiceConfiguration](#)
- [Summary](#)
- [Tag](#)
- [UtilizationMetric](#)
- [UtilizationPreference](#)
- [VolumeConfiguration](#)
- [VolumeRecommendation](#)
- [VolumeRecommendationOption](#)

AccountEnrollmentStatus

Describes the enrollment status of an organization's member accounts in AWS Compute Optimizer.

Contents

accountId

The AWS account ID.

Type: String

Required: No

lastUpdatedTimestamp

The Unix epoch timestamp, in seconds, of when the account enrollment status was last updated.

Type: Timestamp

Required: No

status

The account enrollment status.

Type: String

Valid Values: Active | Inactive | Pending | Failed

Required: No

statusReason

The reason for the account enrollment status.

For example, an account might show a status of Pending because member accounts of an organization require more time to be enrolled in the service.

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

AutoScalingGroupConfiguration

Describes the configuration of an EC2 Auto Scaling group.

Contents

allocationStrategy

Describes the allocation strategy that the EC2 Auto Scaling group uses. This field is only available for EC2 Auto Scaling groups with mixed instance types.

Type: String

Valid Values: `Prioritized` | `LowestPrice`

Required: No

desiredCapacity

The desired capacity, or number of instances, for the EC2 Auto Scaling group.

Type: Integer

Required: No

estimatedInstanceHourReductionPercentage

Describes the projected percentage reduction in instance hours after adopting the recommended configuration. This field is only available for EC2 Auto Scaling groups with scaling policies.

Type: Double

Required: No

instanceType

The instance type for the EC2 Auto Scaling group.

Type: String

Required: No

maxSize

The maximum size, or maximum number of instances, for the EC2 Auto Scaling group.

Type: Integer

Required: No

minSize

The minimum size, or minimum number of instances, for the EC2 Auto Scaling group.

Type: Integer

Required: No

mixedInstanceTypes

List the instance types within an EC2 Auto Scaling group that has mixed instance types.

Type: Array of strings

Required: No

type

Describes whether the EC2 Auto Scaling group has a single instance type or a mixed instance type configuration.

Type: String

Valid Values: `SingleInstanceType` | `MixedInstanceTypes`

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

AutoScalingGroupEstimatedMonthlySavings

An object that describes the estimated monthly savings possible by adopting Compute Optimizer's Auto Scaling group recommendations. This is based on the Savings Plans and Reserved Instances discounts.

Contents

currency

The currency of the estimated monthly savings.

Type: String

Valid Values: USD | CNY

Required: No

value

The value of the estimated monthly savings.

Type: Double

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

AutoScalingGroupRecommendation

Describes an Auto Scaling group recommendation.

Contents

accountId

The AWS account ID of the Auto Scaling group.

Type: String

Required: No

autoScalingGroupArn

The Amazon Resource Name (ARN) of the Auto Scaling group.

Type: String

Required: No

autoScalingGroupName

The name of the Auto Scaling group.

Type: String

Required: No

currentConfiguration

An array of objects that describe the current configuration of the Auto Scaling group.

Type: [AutoScalingGroupConfiguration](#) object

Required: No

currentInstanceGpuInfo

Describes the GPU accelerator settings for the current instance type of the Auto Scaling group.

Type: [GpuInfo](#) object

Required: No

currentPerformanceRisk

The risk of the current Auto Scaling group not meeting the performance needs of its workloads. The higher the risk, the more likely the current Auto Scaling group configuration has insufficient capacity and cannot meet workload requirements.

Type: String

Valid Values: VeryLow | Low | Medium | High

Required: No

effectiveRecommendationPreferences

An object that describes the effective recommendation preferences for the Auto Scaling group.

Type: [EffectiveRecommendationPreferences](#) object

Required: No

finding

The finding classification of the Auto Scaling group.

Findings for Auto Scaling groups include:

- **NotOptimized** —An Auto Scaling group is considered not optimized when AWS Compute Optimizer identifies a recommendation that can provide better performance for your workload.
- **Optimized** —An Auto Scaling group is considered optimized when Compute Optimizer determines that the group is correctly provisioned to run your workload based on the chosen instance type. For optimized resources, Compute Optimizer might recommend a new generation instance type.

Type: String

Valid Values: Underprovisioned | Overprovisioned | Optimized | NotOptimized

Required: No

inferredWorkloadTypes

The applications that might be running on the instances in the Auto Scaling group as inferred by Compute Optimizer.

Compute Optimizer can infer if one of the following applications might be running on the instances:

- `AmazonEmr` - Infers that Amazon EMR might be running on the instances.
- `ApacheCassandra` - Infers that Apache Cassandra might be running on the instances.
- `ApacheHadoop` - Infers that Apache Hadoop might be running on the instances.
- `Memcached` - Infers that Memcached might be running on the instances.
- `NGINX` - Infers that NGINX might be running on the instances.
- `PostgreSql` - Infers that PostgreSQL might be running on the instances.
- `Redis` - Infers that Redis might be running on the instances.
- `Kafka` - Infers that Kafka might be running on the instance.
- `SQLServer` - Infers that SQLServer might be running on the instance.

Type: Array of strings

Valid Values: `AmazonEmr` | `ApacheCassandra` | `ApacheHadoop` | `Memcached` | `Nginx` | `PostgreSql` | `Redis` | `Kafka` | `SQLServer`

Required: No

lastRefreshTimestamp

The timestamp of when the Auto Scaling group recommendation was last generated.

Type: Timestamp

Required: No

lookBackPeriodInDays

The number of days for which utilization metrics were analyzed for the Auto Scaling group.

Type: Double

Required: No

recommendationOptions

An array of objects that describe the recommendation options for the Auto Scaling group.

Type: Array of [AutoScalingGroupRecommendationOption](#) objects

Required: No

utilizationMetrics

An array of objects that describe the utilization metrics of the Auto Scaling group.

Type: Array of [UtilizationMetric](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

AutoScalingGroupRecommendationOption

Describes a recommendation option for an Auto Scaling group.

Contents

configuration

An array of objects that describe an Auto Scaling group configuration.

Type: [AutoScalingGroupConfiguration](#) object

Required: No

instanceGpuInfo

Describes the GPU accelerator settings for the recommended instance type of the Auto Scaling group.

Type: [GpuInfo](#) object

Required: No

migrationEffort

The level of effort required to migrate from the current instance type to the recommended instance type.

For example, the migration effort is Low if Amazon EMR is the inferred workload type and an AWS Graviton instance type is recommended. The migration effort is Medium if a workload type couldn't be inferred but an AWS Graviton instance type is recommended. The migration effort is VeryLow if both the current and recommended instance types are of the same CPU architecture.

Type: String

Valid Values: VeryLow | Low | Medium | High

Required: No

performanceRisk

The performance risk of the Auto Scaling group configuration recommendation.

Performance risk indicates the likelihood of the recommended instance type not meeting the resource needs of your workload. Compute Optimizer calculates an individual performance risk score for each specification of the recommended instance, including CPU, memory, EBS throughput, EBS IOPS, disk throughput, disk IOPS, network throughput, and network PPS. The performance risk of the recommended instance is calculated as the maximum performance risk score across the analyzed resource specifications.

The value ranges from 0 - 4, with 0 meaning that the recommended resource is predicted to always provide enough hardware capability. The higher the performance risk is, the more likely you should validate whether the recommendation will meet the performance requirements of your workload before migrating your resource.

Type: Double

Valid Range: Minimum value of 0. Maximum value of 4.

Required: No

projectedUtilizationMetrics

An array of objects that describe the projected utilization metrics of the Auto Scaling group recommendation option.

Note

The Cpu and Memory metrics are the only projected utilization metrics returned. Additionally, the Memory metric is returned only for resources that have the unified CloudWatch agent installed on them. For more information, see [Enabling Memory Utilization with the CloudWatch Agent](#).

Type: Array of [UtilizationMetric](#) objects

Required: No

rank

The rank of the Auto Scaling group recommendation option.

The top recommendation option is ranked as 1.

Type: Integer

Required: No

savingsOpportunity

An object that describes the savings opportunity for the Auto Scaling group recommendation option. Savings opportunity includes the estimated monthly savings amount and percentage.

Type: [SavingsOpportunity](#) object

Required: No

savingsOpportunityAfterDiscounts

An object that describes the savings opportunity for the Auto Scaling group recommendation option that includes Savings Plans and Reserved Instances discounts. Savings opportunity includes the estimated monthly savings and percentage.

Type: [AutoScalingGroupSavingsOpportunityAfterDiscounts](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

AutoScalingGroupSavingsOpportunityAfterDiscounts

Describes the savings opportunity for Auto Scaling group recommendations after applying the Savings Plans and Reserved Instances discounts.

Savings opportunity represents the estimated monthly savings you can achieve by implementing Compute Optimizer recommendations.

Contents

estimatedMonthlySavings

An object that describes the estimated monthly savings possible by adopting Compute Optimizer's Auto Scaling group recommendations. This is based on the Savings Plans and Reserved Instances pricing discounts.

Type: [AutoScalingGroupEstimatedMonthlySavings](#) object

Required: No

savingsOpportunityPercentage

The estimated monthly savings possible as a percentage of monthly cost after applying the Savings Plans and Reserved Instances discounts. This saving can be achieved by adopting Compute Optimizer's Auto Scaling group recommendations.

Type: Double

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

ContainerConfiguration

Describes the container configurations within the tasks of your Amazon ECS service.

Contents

containerName

The name of the container.

Type: String

Required: No

cpu

The number of CPU units reserved for the container.

Type: Integer

Required: No

memorySizeConfiguration

The memory size configurations for the container.

Type: [MemorySizeConfiguration](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

ContainerRecommendation

The CPU and memory recommendations for a container within the tasks of your Amazon ECS service.

Contents

containerName

The name of the container.

Type: String

Required: No

cpu

The recommended number of CPU units reserved for the container.

Type: Integer

Required: No

memorySizeConfiguration

The recommended memory size configurations for the container.

Type: [MemorySizeConfiguration](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

CurrentPerformanceRiskRatings

Describes the performance risk ratings for a given resource type.

Resources with a high or medium rating are at risk of not meeting the performance needs of their workloads, while resources with a low rating are performing well in their workloads.

Contents

high

A count of the applicable resource types with a high performance risk rating.

Type: Long

Required: No

low

A count of the applicable resource types with a low performance risk rating.

Type: Long

Required: No

medium

A count of the applicable resource types with a medium performance risk rating.

Type: Long

Required: No

veryLow

A count of the applicable resource types with a very low performance risk rating.

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

CustomizableMetricParameters

Defines the various metric parameters that can be customized, such as threshold and headroom.

Contents

headroom

The headroom value in percentage used for the specified metric parameter.

The following lists the valid values for CPU and memory utilization.

- CPU utilization: PERCENT_30 | PERCENT_20 | PERCENT_0
- Memory utilization: PERCENT_30 | PERCENT_20 | PERCENT_10

Type: String

Valid Values: PERCENT_30 | PERCENT_20 | PERCENT_10 | PERCENT_0

Required: No

threshold

The threshold value used for the specified metric parameter.

Note

You can only specify the threshold value for CPU utilization.

Type: String

Valid Values: P90 | P95 | P99_5

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

DBStorageConfiguration

The configuration of the recommended RDS storage.

Contents

allocatedStorage

The size of the DB storage in gigabytes (GB).

Type: Integer

Required: No

iops

The provisioned IOPs of the DB storage.

Type: Integer

Required: No

maxAllocatedStorage

The maximum limit in gibibytes (GiB) to which Amazon RDS can automatically scale the storage of the DB instance.

Type: Integer

Required: No

storageThroughput

The storage throughput of the DB storage.

Type: Integer

Required: No

storageType

The type of DB storage.

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

EBSEffectiveRecommendationPreferences

Describes the effective recommendation preferences for Amazon EBS volumes.

Contents

savingsEstimationMode

Describes the savings estimation mode preference applied for calculating savings opportunity for Amazon EBS volumes.

Type: [EBSSavingsEstimationMode](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

EBSEstimatedMonthlySavings

An object that describes the estimated monthly savings possible by adopting Compute Optimizer's Amazon EBS volume recommendations. This includes any applicable discounts.

Contents

currency

The currency of the estimated monthly savings.

Type: String

Valid Values: USD | CNY

Required: No

value

The value of the estimated monthly savings.

Type: Double

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

EBSFilter

Describes a filter that returns a more specific list of Amazon Elastic Block Store (Amazon EBS) volume recommendations. Use this filter with the [GetEBSVolumeRecommendations](#) action.

You can use `LambdaFunctionRecommendationFilter` with the [GetLambdaFunctionRecommendations](#) action, `JobFilter` with the [DescribeRecommendationExportJobs](#) action, and `Filter` with the [GetAutoScalingGroupRecommendations](#) and [GetEC2InstanceRecommendations](#) actions.

Contents

name

The name of the filter.

Specify `Finding` to return recommendations with a specific finding classification (for example, `NotOptimized`).

You can filter your Amazon EBS volume recommendations by `tag:key` and `tag-key` tags.

A `tag:key` is a key and value combination of a tag assigned to your Amazon EBS volume recommendations. Use the tag key in the filter name and the tag value as the filter value. For example, to find all Amazon EBS volume recommendations that have a tag with the key of `Owner` and the value of `TeamA`, specify `tag:Owner` for the filter name and `TeamA` for the filter value.

A `tag-key` is the key of a tag assigned to your Amazon EBS volume recommendations. Use this filter to find all of your Amazon EBS volume recommendations that have a tag with a specific key. This doesn't consider the tag value. For example, you can find your Amazon EBS volume recommendations with a tag key value of `Owner` or without any tag keys assigned.

Type: String

Valid Values: `Finding`

Required: No

values

The value of the filter.

The valid values are `Optimized`, or `NotOptimized`.

Type: Array of strings

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

EBSSavingsEstimationMode

Describes the savings estimation mode used for calculating savings opportunity for Amazon EBS volumes.

Contents

source

Describes the source for calculating the savings opportunity for Amazon EBS volumes.

Type: String

Valid Values: `PublicPricing` | `CostExplorerRightsizing` | `CostOptimizationHub`

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

EBSSavingsOpportunityAfterDiscounts

Describes the savings opportunity for Amazon EBS volume recommendations after applying specific discounts.

Contents

estimatedMonthlySavings

The estimated monthly savings possible as a percentage of monthly cost by adopting Compute Optimizer's Amazon EBS volume recommendations. This saving includes any applicable discounts.

Type: [EBSEstimatedMonthlySavings](#) object

Required: No

savingsOpportunityPercentage

The estimated monthly savings possible as a percentage of monthly cost after applying the specific discounts. This saving can be achieved by adopting Compute Optimizer's Amazon EBS volume recommendations.

Type: Double

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

EBSUtilizationMetric

Describes a utilization metric of an Amazon Elastic Block Store (Amazon EBS) volume.

Compare the utilization metric data of your resource against its projected utilization metric data to determine the performance difference between your current resource and the recommended option.

Contents

name

The name of the utilization metric.

The following utilization metrics are available:

- `VolumeReadOpsPerSecond` - The completed read operations per second from the volume in a specified period of time.

Unit: Count

- `VolumeWriteOpsPerSecond` - The completed write operations per second to the volume in a specified period of time.

Unit: Count

- `VolumeReadBytesPerSecond` - The bytes read per second from the volume in a specified period of time.

Unit: Bytes

- `VolumeWriteBytesPerSecond` - The bytes written to the volume in a specified period of time.

Unit: Bytes

Type: String

Valid Values: `VolumeReadOpsPerSecond` | `VolumeWriteOpsPerSecond` | `VolumeReadBytesPerSecond` | `VolumeWriteBytesPerSecond`

Required: No

statistic

The statistic of the utilization metric.

The Compute Optimizer API, AWS Command Line Interface (AWS CLI), and SDKs return utilization metrics using only the `Maximum` statistic, which is the highest value observed during the specified period.

The Compute Optimizer console displays graphs for some utilization metrics using the `Average` statistic, which is the value of `Sum / SampleCount` during the specified period. For more information, see [Viewing resource recommendations](#) in the *AWS Compute Optimizer User Guide*. You can also get averaged utilization metric data for your resources using Amazon CloudWatch. For more information, see the [Amazon CloudWatch User Guide](#).

Type: String

Valid Values: `Maximum` | `Average`

Required: No

value

The value of the utilization metric.

Type: Double

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

ECSEffectiveRecommendationPreferences

Describes the effective recommendation preferences for Amazon ECS services.

Contents

savingsEstimationMode

Describes the savings estimation mode preference applied for calculating savings opportunity for Amazon ECS services.

Type: [ECSSavingsEstimationMode](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

ECSEstimatedMonthlySavings

Describes the estimated monthly savings possible for Amazon ECS services by adopting Compute Optimizer recommendations. This is based on Amazon ECS service pricing after applying Savings Plans discounts.

Contents

currency

The currency of the estimated monthly savings.

Type: String

Valid Values: USD | CNY

Required: No

value

The value of the estimated monthly savings for Amazon ECS services.

Type: Double

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

ECSSavingsEstimationMode

Describes the savings estimation mode used for calculating savings opportunity for Amazon ECS services.

Contents

source

Describes the source for calculating the savings opportunity for Amazon ECS services.

Type: String

Valid Values: `PublicPricing` | `CostExplorerRightsizing` | `CostOptimizationHub`

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

ECSSavingsOpportunityAfterDiscounts

Describes the savings opportunity for Amazon ECS service recommendations after applying Savings Plans discounts.

Savings opportunity represents the estimated monthly savings after applying Savings Plans discounts. You can achieve this by implementing a given Compute Optimizer recommendation.

Contents

estimatedMonthlySavings

The estimated monthly savings possible by adopting Compute Optimizer's Amazon ECS service recommendations. This includes any applicable Savings Plans discounts.

Type: [ECSEstimatedMonthlySavings](#) object

Required: No

savingsOpportunityPercentage

The estimated monthly savings possible as a percentage of monthly cost by adopting Compute Optimizer's Amazon ECS service recommendations. This includes any applicable Savings Plans discounts.

Type: Double

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

ECSServiceProjectedMetric

Describes the projected metrics of an Amazon ECS service recommendation option.

To determine the performance difference between your current Amazon ECS service and the recommended option, compare the metric data of your service against its projected metric data.

Contents

lowerBoundValues

The lower bound values for the projected metric.

Type: Array of doubles

Required: No

name

The name of the projected metric.

The following metrics are available:

- **Cpu** — The percentage of allocated compute units that are currently in use on the service tasks.
- **Memory** — The percentage of memory that's currently in use on the service tasks.

Type: String

Valid Values: Cpu | Memory

Required: No

timestamps

The timestamps of the projected metric.

Type: Array of timestamps

Required: No

upperBoundValues

The upper bound values for the projected metric.

Type: Array of doubles

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

ECSServiceProjectedUtilizationMetric

Describes the projected utilization metrics of an Amazon ECS service recommendation option.

To determine the performance difference between your current Amazon ECS service and the recommended option, compare the utilization metric data of your service against its projected utilization metric data.

Contents

lowerBoundValue

The lower bound values for the projected utilization metrics.

Type: Double

Required: No

name

The name of the projected utilization metric.

The following utilization metrics are available:

- **Cpu** — The percentage of allocated compute units that are currently in use on the service tasks.
- **Memory** — The percentage of memory that's currently in use on the service tasks.

Type: String

Valid Values: Cpu | Memory

Required: No

statistic

The statistic of the projected utilization metric.

The Compute Optimizer API, AWS Command Line Interface (AWS CLI), and SDKs return utilization metrics using only the **Maximum** statistic, which is the highest value observed during the specified period.

The Compute Optimizer console displays graphs for some utilization metrics using the **Average** statistic, which is the value of `Sum / SampleCount` during the specified period. For more

information, see [Viewing resource recommendations](#) in the *AWS Compute Optimizer User Guide*. You can also get averaged utilization metric data for your resources using Amazon CloudWatch. For more information, see the [Amazon CloudWatch User Guide](#).

Type: String

Valid Values: Maximum | Average

Required: No

upperBoundValue

The upper bound values for the projected utilization metrics.

Type: Double

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

ECSServiceRecommendation

Describes an Amazon ECS service recommendation.

Contents

accountId

The AWS account ID of the Amazon ECS service.

Type: String

Required: No

currentPerformanceRisk

The risk of the current Amazon ECS service not meeting the performance needs of its workloads. The higher the risk, the more likely the current service can't meet the performance requirements of its workload.

Type: String

Valid Values: VeryLow | Low | Medium | High

Required: No

currentServiceConfiguration

The configuration of the current Amazon ECS service.

Type: [ServiceConfiguration](#) object

Required: No

effectiveRecommendationPreferences

Describes the effective recommendation preferences for Amazon ECS services.

Type: [ECSEffectiveRecommendationPreferences](#) object

Required: No

finding

The finding classification of an Amazon ECS service.

Findings for Amazon ECS services include:

- **Underprovisioned** — When Compute Optimizer detects that there's not enough memory or CPU, an Amazon ECS service is considered under-provisioned. An under-provisioned service might result in poor application performance.
- **Overprovisioned** — When Compute Optimizer detects that there's excessive memory or CPU, an Amazon ECS service is considered over-provisioned. An over-provisioned service might result in additional infrastructure costs.
- **Optimized** — When both the CPU and memory of your Amazon ECS service meet the performance requirements of your workload, the service is considered optimized.

Type: String

Valid Values: Optimized | Underprovisioned | Overprovisioned

Required: No

findingReasonCodes

The reason for the finding classification of an Amazon ECS service.

Finding reason codes for Amazon ECS services include:

- **CPUUnderprovisioned** — The service CPU configuration can be sized up to enhance the performance of your workload. This is identified by analyzing the CPUUtilization metric of the current service during the look-back period.
- **CPUOverprovisioned** — The service CPU configuration can be sized down while still meeting the performance requirements of your workload. This is identified by analyzing the CPUUtilization metric of the current service during the look-back period.
- **MemoryUnderprovisioned** — The service memory configuration can be sized up to enhance the performance of your workload. This is identified by analyzing the MemoryUtilization metric of the current service during the look-back period.
- **MemoryOverprovisioned** — The service memory configuration can be sized down while still meeting the performance requirements of your workload. This is identified by analyzing the MemoryUtilization metric of the current service during the look-back period.

Type: Array of strings

Valid Values: MemoryOverprovisioned | MemoryUnderprovisioned | CPUOverprovisioned | CPUUnderprovisioned

Required: No

lastRefreshTimestamp


The timestamp of when the Amazon ECS service recommendation was last generated.

Type: Timestamp

Required: No

launchType

The launch type the Amazon ECS service is using.



Note

Compute Optimizer only supports the Fargate launch type.

Type: String

Valid Values: EC2 | Fargate

Required: No

lookbackPeriodInDays

The number of days the Amazon ECS service utilization metrics were analyzed.

Type: Double

Required: No

serviceArn

The Amazon Resource Name (ARN) of the current Amazon ECS service.

The following is the format of the ARN:

```
arn:aws:ecs:region:aws_account_id:service/cluster-name/service-name
```

Type: String

Required: No

serviceRecommendationOptions

An array of objects that describe the recommendation options for the Amazon ECS service.

Type: Array of [ECSServiceRecommendationOption](#) objects

Required: No

tags

A list of tags assigned to your Amazon ECS service recommendations.

Type: Array of [Tag](#) objects

Required: No

utilizationMetrics

An array of objects that describe the utilization metrics of the Amazon ECS service.

Type: Array of [ECSServiceUtilizationMetric](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

ECSServiceRecommendationFilter

Describes a filter that returns a more specific list of Amazon ECS service recommendations. Use this filter with the [GetECSServiceRecommendations](#) action.

Contents

name

The name of the filter.

Specify `Finding` to return recommendations with a specific finding classification.

Specify `FindingReasonCode` to return recommendations with a specific finding reason code.

You can filter your Amazon ECS service recommendations by `tag:key` and `tag-key` tags.

A `tag:key` is a key and value combination of a tag assigned to your Amazon ECS service recommendations. Use the tag key in the filter name and the tag value as the filter value. For example, to find all Amazon ECS service recommendations that have a tag with the key of `Owner` and the value of `TeamA`, specify `tag:Owner` for the filter name and `TeamA` for the filter value.

A `tag-key` is the key of a tag assigned to your Amazon ECS service recommendations. Use this filter to find all of your Amazon ECS service recommendations that have a tag with a specific key. This doesn't consider the tag value. For example, you can find your Amazon ECS service recommendations with a tag key value of `Owner` or without any tag keys assigned.

Type: String

Valid Values: `Finding` | `FindingReasonCode`

Required: No

values

The value of the filter.

The valid values for this parameter are as follows:

- If you specify the name parameter as `Finding`, specify `Optimized`, `Underprovisioned`, or `Overprovisioned`.

- If you specify the name parameter as `FindingReasonCode`, specify `CPUUnderprovisioned`, `CPUOverprovisioned`, `MemoryUnderprovisioned`, or `MemoryOverprovisioned`.

Type: Array of strings

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

ECSServiceRecommendationOption

Describes the recommendation options for an Amazon ECS service.

Contents

containerRecommendations

The CPU and memory size recommendations for the containers within the task of your Amazon ECS service.

Type: Array of [ContainerRecommendation](#) objects

Required: No

cpu

The CPU size of the Amazon ECS service recommendation option.

Type: Integer

Required: No

memory

The memory size of the Amazon ECS service recommendation option.

Type: Integer

Required: No

projectedUtilizationMetrics

An array of objects that describe the projected utilization metrics of the Amazon ECS service recommendation option.

Type: Array of [ECSServiceProjectedUtilizationMetric](#) objects

Required: No

savingsOpportunity

Describes the savings opportunity for recommendations of a given resource type or for the recommendation option of an individual resource.

Savings opportunity represents the estimated monthly savings you can achieve by implementing a given Compute Optimizer recommendation.

⚠ Important

Savings opportunity data requires that you opt in to Cost Explorer, as well as activate **Receive Amazon EC2 resource recommendations** in the Cost Explorer preferences page. That creates a connection between Cost Explorer and Compute Optimizer. With this connection, Cost Explorer generates savings estimates considering the price of existing resources, the price of recommended resources, and historical usage data. Estimated monthly savings reflects the projected dollar savings associated with each of the recommendations generated. For more information, see [Enabling Cost Explorer](#) and [Optimizing your cost with Rightsizing Recommendations](#) in the *Cost Management User Guide*.

Type: [SavingsOpportunity](#) object

Required: No

savingsOpportunityAfterDiscounts

Describes the savings opportunity for Amazon ECS service recommendations or for the recommendation option.

Savings opportunity represents the estimated monthly savings after applying Savings Plans discounts. You can achieve this by implementing a given Compute Optimizer recommendation.

Type: [ECSSavingsOpportunityAfterDiscounts](#) 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 Java V2](#)

- [AWS SDK for Ruby V3](#)

ECSServiceRecommendedOptionProjectedMetric

Describes the projected metrics of an Amazon ECS service recommendation option.

To determine the performance difference between your current Amazon ECS service and the recommended option, compare the metric data of your service against its projected metric data.

Contents

projectedMetrics

An array of objects that describe the projected metric.

Type: Array of [ECSServiceProjectedMetric](#) objects

Required: No

recommendedCpuUnits

The recommended CPU size for the Amazon ECS service.

Type: Integer

Required: No

recommendedMemorySize

The recommended memory size for the Amazon ECS service.

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

ECSServiceUtilizationMetric

Describes the utilization metric of an Amazon ECS service.

To determine the performance difference between your current Amazon ECS service and the recommended option, compare the utilization metric data of your service against its projected utilization metric data.

Contents

name

The name of the utilization metric.

The following utilization metrics are available:

- `Cpu` — The amount of CPU capacity that's used in the service.
- `Memory` — The amount of memory that's used in the service.

Type: String

Valid Values: `Cpu` | `Memory`

Required: No

statistic

The statistic of the utilization metric.

The Compute Optimizer API, AWS Command Line Interface (AWS CLI), and SDKs return utilization metrics using only the `Maximum` statistic, which is the highest value observed during the specified period.

The Compute Optimizer console displays graphs for some utilization metrics using the `Average` statistic, which is the value of `Sum / SampleCount` during the specified period. For more information, see [Viewing resource recommendations](#) in the *AWS Compute Optimizer User Guide*. You can also get averaged utilization metric data for your resources using Amazon CloudWatch. For more information, see the [Amazon CloudWatch User Guide](#).

Type: String

Valid Values: `Maximum` | `Average`

Required: No

value

The value of the utilization metric.

Type: Double

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

EffectivePreferredResource

Describes the effective preferred resources that Compute Optimizer considers as rightsizing recommendation candidates.

Note

Compute Optimizer only supports Amazon EC2 instance types.

Contents

effectiveIncludeList

The expanded version of your preferred resource's include list.

Type: Array of strings

Required: No

excludeList

The list of preferred resources values that you want excluded from rightsizing recommendation candidates.

Type: Array of strings

Required: No

includeList

The list of preferred resource values that you want considered as rightsizing recommendation candidates.

Type: Array of strings

Required: No

name

The name of the preferred resource list.

Type: String

Valid Values: `Ec2InstanceTypes`

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

EffectiveRecommendationPreferences

Describes the effective recommendation preferences for a resource.

Contents

cpuVendorArchitectures

Describes the CPU vendor and architecture for an instance or Auto Scaling group recommendations.

For example, when you specify `AWS_ARM64` with:

- A [GetEC2InstanceRecommendations](#) or [GetAutoScalingGroupRecommendations](#) request, Compute Optimizer returns recommendations that consist of Graviton instance types only.
- A [GetEC2RecommendationProjectedMetrics](#) request, Compute Optimizer returns projected utilization metrics for Graviton instance type recommendations only.
- A [ExportEC2InstanceRecommendations](#) or [ExportAutoScalingGroupRecommendations](#) request, Compute Optimizer exports recommendations that consist of Graviton instance types only.

Type: Array of strings

Valid Values: `AWS_ARM64` | `CURRENT`

Required: No

enhancedInfrastructureMetrics

Describes the activation status of the enhanced infrastructure metrics preference.

A status of `Active` confirms that the preference is applied in the latest recommendation refresh, and a status of `Inactive` confirms that it's not yet applied to recommendations.

For more information, see [Enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Type: String

Valid Values: `Active` | `Inactive`

Required: No

externalMetricsPreference

An object that describes the external metrics recommendation preference.

If the preference is applied in the latest recommendation refresh, an object with a valid `source` value appears in the response. If the preference isn't applied to the recommendations already, then this object doesn't appear in the response.

Type: [ExternalMetricsPreference](#) object

Required: No

inferredWorkloadTypes

Describes the activation status of the inferred workload types preference.

A status of `Active` confirms that the preference is applied in the latest recommendation refresh. A status of `Inactive` confirms that it's not yet applied to recommendations.

Type: String

Valid Values: `Active` | `Inactive`

Required: No

lookBackPeriod

The number of days the utilization metrics of the AWS resource are analyzed.

Type: String

Valid Values: `DAYS_14` | `DAYS_32` | `DAYS_93`

Required: No

preferredResources

The resource type values that are considered as candidates when generating rightsizing recommendations.

Type: Array of [EffectivePreferredResource](#) objects

Required: No

savingsEstimationMode

Describes the savings estimation mode applied for calculating savings opportunity for a resource.

Type: [InstanceSavingsEstimationMode](#) object

Required: No

utilizationPreferences

The resource's CPU and memory utilization preferences, such as threshold and headroom, that are used to generate rightsizing recommendations.

Note

This preference is only available for the Amazon EC2 instance resource type.

Type: Array of [UtilizationPreference](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

EnrollmentFilter

Describes a filter that returns a more specific list of account enrollment statuses. Use this filter with the [GetEnrollmentStatusesForOrganization](#) action.

Contents

name

The name of the filter.

Specify Status to return accounts with a specific enrollment status (for example, Active).

Type: String

Valid Values: Status

Required: No

values

The value of the filter.

The valid values are Active, Inactive, Pending, and Failed.

Type: Array of strings

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

EstimatedMonthlySavings

Describes the estimated monthly savings amount possible, based on On-Demand instance pricing, by adopting Compute Optimizer recommendations for a given resource.

For more information, see [Estimated monthly savings and savings opportunities](#) in the *AWS Compute Optimizer User Guide*.

Contents

currency

The currency of the estimated monthly savings.

Type: String

Valid Values: USD

Required: No

value

The value of the estimated monthly savings.

Type: Double

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

ExportDestination

Describes the destination of the recommendations export and metadata files.

Contents

s3

An object that describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and object keys of a recommendations export file, and its associated metadata file.

Type: [S3Destination](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

ExternalMetricsPreference

Describes the external metrics preferences for EC2 rightsizing recommendations.

Contents

source

Contains the source options for external metrics preferences.

Type: String

Valid Values: Datadog | Dynatrace | NewRelic | Instana

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

ExternalMetricStatus

Describes Compute Optimizer's integration status with your chosen external metric provider. For example, Datadog.

Contents

statusCode

The status code for Compute Optimizer's integration with an external metrics provider.

Type: String

Valid Values: NO_EXTERNAL_METRIC_SET | INTEGRATION_SUCCESS | DATADOG_INTEGRATION_ERROR | DYNATRACE_INTEGRATION_ERROR | NEWRELIC_INTEGRATION_ERROR | INSTANA_INTEGRATION_ERROR | INSUFFICIENT_DATADOG_METRICS | INSUFFICIENT_DYNATRACE_METRICS | INSUFFICIENT_NEWRELIC_METRICS | INSUFFICIENT_INSTANA_METRICS

Required: No

statusReason

The reason for Compute Optimizer's integration status with your external metric provider.

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

Filter

Describes a filter that returns a more specific list of recommendations. Use this filter with the [GetAutoScalingGroupRecommendations](#) and [GetEC2InstanceRecommendations](#) actions.

You can use `EBSFilter` with the [GetEBSVolumeRecommendations](#) action, `LambdaFunctionRecommendationFilter` with the [GetLambdaFunctionRecommendations](#) action, and `JobFilter` with the [DescribeRecommendationExportJobs](#) action.

Contents

name

The name of the filter.

Specify `Finding` to return recommendations with a specific finding classification. For example, `Underprovisioned`.

Specify `RecommendationSourceType` to return recommendations of a specific resource type. For example, `Ec2Instance`.

Specify `FindingReasonCodes` to return recommendations with a specific finding reason code. For example, `CPUUnderprovisioned`.

Specify `InferredWorkloadTypes` to return recommendations of a specific inferred workload. For example, `Redis`.

You can filter your EC2 instance recommendations by `tag:key` and `tag-key` tags.

A `tag:key` is a key and value combination of a tag assigned to your recommendations. Use the tag key in the filter name and the tag value as the filter value. For example, to find all recommendations that have a tag with the key of `Owner` and the value of `TeamA`, specify `tag:Owner` for the filter name and `TeamA` for the filter value.

A `tag-key` is the key of a tag assigned to your recommendations. Use this filter to find all of your recommendations that have a tag with a specific key. This doesn't consider the tag value. For example, you can find your recommendations with a tag key value of `Owner` or without any tag keys assigned.

Type: String

Valid Values: Finding | FindingReasonCodes | RecommendationSourceType | InferredWorkloadTypes

Required: No

values

The value of the filter.

The valid values for this parameter are as follows, depending on what you specify for the name parameter and the resource type that you wish to filter results for:

- Specify `Optimized` or `NotOptimized` if you specify the name parameter as `Finding` and you want to filter results for Auto Scaling groups.
- Specify `Underprovisioned`, `Overprovisioned`, or `Optimized` if you specify the name parameter as `Finding` and you want to filter results for EC2 instances.
- Specify `Ec2Instance` or `AutoScalingGroup` if you specify the name parameter as `RecommendationSourceType`.
- Specify one of the following options if you specify the name parameter as `FindingReasonCodes`:
 - **CPUOverprovisioned** — The instance's CPU configuration can be sized down while still meeting the performance requirements of your workload.
 - **CPUUnderprovisioned** — The instance's CPU configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better CPU performance.
 - **MemoryOverprovisioned** — The instance's memory configuration can be sized down while still meeting the performance requirements of your workload.
 - **MemoryUnderprovisioned** — The instance's memory configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better memory performance.
 - **EBSThroughputOverprovisioned** — The instance's EBS throughput configuration can be sized down while still meeting the performance requirements of your workload.
 - **EBSThroughputUnderprovisioned** — The instance's EBS throughput configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better EBS throughput performance.
 - **EBSIOPSOverprovisioned** — The instance's EBS IOPS configuration can be sized down while still meeting the performance requirements of your workload.

- **EBSIOPSUnderprovisioned** — The instance's EBS IOPS configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better EBS IOPS performance.
- **NetworkBandwidthOverprovisioned** — The instance's network bandwidth configuration can be sized down while still meeting the performance requirements of your workload.
- **NetworkBandwidthUnderprovisioned** — The instance's network bandwidth configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better network bandwidth performance. This finding reason happens when the NetworkIn or NetworkOut performance of an instance is impacted.
- **NetworkPPSOverprovisioned** — The instance's network PPS (packets per second) configuration can be sized down while still meeting the performance requirements of your workload.
- **NetworkPPSUnderprovisioned** — The instance's network PPS (packets per second) configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better network PPS performance.
- **DiskIOPSOverprovisioned** — The instance's disk IOPS configuration can be sized down while still meeting the performance requirements of your workload.
- **DiskIOPSUnderprovisioned** — The instance's disk IOPS configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better disk IOPS performance.
- **DiskThroughputOverprovisioned** — The instance's disk throughput configuration can be sized down while still meeting the performance requirements of your workload.
- **DiskThroughputUnderprovisioned** — The instance's disk throughput configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better disk throughput performance.

Type: Array of strings

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

GetRecommendationError

Describes an error experienced when getting recommendations.

For example, an error is returned if you request recommendations for an unsupported Auto Scaling group, or if you request recommendations for an instance of an unsupported instance family.

Contents

code

The error code.

Type: String

Required: No

identifier

The ID of the error.

Type: String

Required: No

message

The message, or reason, for the error.

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

Gpu

Describes the GPU accelerators for the instance type.

Contents

gpuCount

The number of GPUs for the instance type.

Type: Integer

Required: No

gpuMemorySizeInMiB

The total size of the memory for the GPU accelerators for the instance type, in MiB.

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

GpuInfo

Describes the GPU accelerator settings for the instance type.

Contents

gpus

Describes the GPU accelerators for the instance type.

Type: Array of [Gpu](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

IdleEstimatedMonthlySavings

Describes the estimated monthly savings possible for idle resources by adopting Compute Optimizer recommendations.

Contents

currency

The currency of the estimated monthly savings.

Type: String

Valid Values: USD | CNY

Required: No

value

The value of the estimated monthly savings for Idle resources.

Type: Double

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

IdleRecommendation

Describes an Idle resource recommendation.

Contents

accountId

The AWS account ID of the idle resource.

Type: String

Required: No

finding

The finding classification of an idle resource.

Type: String

Valid Values: Idle | Unattached

Required: No

findingDescription

A summary of the findings for the resource.

Type: String

Required: No

lastRefreshTimestamp

The timestamp of when the idle resource recommendation was last generated.

Type: Timestamp

Required: No

lookBackPeriodInDays

The number of days the idle resource utilization metrics were analyzed.

Type: Double

Required: No

resourceArn

The ARN of the current idle resource.

Type: String

Required: No

resourceId

The unique identifier for the resource.

Type: String

Required: No

resourceType

The type of resource that is idle.

Type: String

Valid Values: EC2Instance | AutoScalingGroup | EBSVolume | ECSService | RDSDBInstance

Required: No

savingsOpportunity

The savings opportunity for the idle resource.

Type: [IdleSavingsOpportunity](#) object

Required: No

savingsOpportunityAfterDiscounts

The savings opportunity for the idle resource after any applying discounts.

Type: [IdleSavingsOpportunityAfterDiscounts](#) object

Required: No

tags

A list of tags assigned to your idle resource recommendations.

Type: Array of [Tag](#) objects

Required: No

utilizationMetrics

An array of objects that describe the utilization metrics of the idle resource.

Type: Array of [IdleUtilizationMetric](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

IdleRecommendationError

Returns of list of resources that doesn't have idle recommendations.

Contents

code

The error code.

Type: String

Required: No

identifier

The ID of the error.

Type: String

Required: No

message

The error message.

Type: String

Required: No

resourceType

The type of resource associated with the error.

Type: String

Valid Values: EC2Instance | AutoScalingGroup | EBSVolume | ECSService |
RDSDBInstance

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

IdleRecommendationFilter

Describes a filter that returns a more specific list of idle resource recommendations.

Contents

name

The name of the filter.

Specify `Finding` to return recommendations with a specific finding classification.

You can filter your idle resource recommendations by `tag:key` and `tag-key` tags.

A `tag:key` is a key and value combination of a tag assigned to your idle resource recommendations. Use the tag key in the filter name and the tag value as the filter value. For example, to find all idle resource service recommendations that have a tag with the key of `Owner` and the value of `TeamA`, specify `tag:Owner` for the filter name and `TeamA` for the filter value.

A `tag-key` is the key of a tag assigned to your idle resource recommendations. Use this filter to find all of your idle resource recommendations that have a tag with a specific key. This doesn't consider the tag value. For example, you can find your idle resource service recommendations with a tag key value of `Owner` or without any tag keys assigned.

Type: String

Valid Values: `Finding` | `ResourceType`

Required: No

values

The value of the filter.

Type: Array of strings

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

IdleSavingsOpportunity

Describes the savings opportunity for idle resource recommendations.

Contents

estimatedMonthlySavings

The estimated monthly savings possible by adopting Compute Optimizer's idle resource recommendations.

Type: [IdleEstimatedMonthlySavings](#) object

Required: No

savingsOpportunityPercentage

The estimated monthly savings possible as a percentage of monthly cost by adopting Compute Optimizer's idle resource recommendations.

Type: Double

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

IdleSavingsOpportunityAfterDiscounts

Describes the savings opportunity for idle resource recommendations after applying discounts.

Savings opportunity represents the estimated monthly savings after applying discounts. You can achieve this by implementing a given Compute Optimizer recommendation.

Contents

estimatedMonthlySavings

The estimated monthly savings possible by adopting Compute Optimizer's idle resource recommendations. This includes any applicable discounts.

Type: [IdleEstimatedMonthlySavings](#) object

Required: No

savingsOpportunityPercentage

The estimated monthly savings possible as a percentage of monthly cost by adopting Compute Optimizer's idle resource recommendations. This includes any applicable discounts.

Type: Double

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

IdleSummary

Describes the findings summary of the idle resources.

Contents

name

The name of the finding group for the idle resources.

Type: String

Valid Values: Idle | Unattached

Required: No

value

The count of idle resources in the finding group.

Type: Double

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

IdleUtilizationMetric

Describes the utilization metric of an idle resource.

Contents

name

The name of the utilization metric.

Type: String

Valid Values: CPU | Memory | NetworkOutBytesPerSecond | NetworkInBytesPerSecond | DatabaseConnections | EBSVolumeReadIOPS | EBSVolumeWriteIOPS | VolumeReadOpsPerSecond | VolumeWriteOpsPerSecond

Required: No

statistic

The statistic of the utilization metric.

The Compute Optimizer API, AWS Command Line Interface (AWS CLI), and SDKs return utilization metrics using only the Maximum statistic, which is the highest value observed during the specified period.

The Compute Optimizer console displays graphs for some utilization metrics using the Average statistic, which is the value of $\text{Sum} / \text{SampleCount}$ during the specified period. For more information, see [Viewing resource recommendations](#) in the *AWS Compute Optimizer User Guide*. You can also get averaged utilization metric data for your resources using Amazon CloudWatch. For more information, see the [Amazon CloudWatch User Guide](#).

Type: String

Valid Values: Maximum | Average

Required: No

value

The value of the utilization metric.

Type: Double

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

InferredWorkloadSaving

The estimated monthly savings after you adjust the configurations of your instances running on the inferred workload types to the recommended configurations. If the `inferredWorkloadTypes` list contains multiple entries, then the savings are the sum of the monthly savings from instances that run the exact combination of the inferred workload types.

Contents

estimatedMonthlySavings

An object that describes the estimated monthly savings amount possible by adopting Compute Optimizer recommendations for a given resource. This is based on the On-Demand instance pricing.

Type: [EstimatedMonthlySavings](#) object

Required: No

inferredWorkloadTypes

The applications that might be running on the instance as inferred by Compute Optimizer.

Compute Optimizer can infer if one of the following applications might be running on the instance:

- `AmazonEmr` - Infers that Amazon EMR might be running on the instance.
- `ApacheCassandra` - Infers that Apache Cassandra might be running on the instance.
- `ApacheHadoop` - Infers that Apache Hadoop might be running on the instance.
- `Memcached` - Infers that Memcached might be running on the instance.
- `NGINX` - Infers that NGINX might be running on the instance.
- `PostgreSql` - Infers that PostgreSQL might be running on the instance.
- `Redis` - Infers that Redis might be running on the instance.
- `Kafka` - Infers that Kafka might be running on the instance.
- `SQLServer` - Infers that SQLServer might be running on the instance.

Type: Array of strings

Valid Values: `AmazonEmr` | `ApacheCassandra` | `ApacheHadoop` | `Memcached` | `Nginx` | `PostgreSql` | `Redis` | `Kafka` | `SQLServer`

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

InstanceEstimatedMonthlySavings

An object that describes the estimated monthly savings possible by adopting Compute Optimizer's Amazon EC2 instance recommendations. This is based on the Savings Plans and Reserved Instances pricing discounts.

Contents

currency

The currency of the estimated monthly savings.

Type: String

Valid Values: USD | CNY

Required: No

value

The value of the estimated monthly savings.

Type: Double

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

InstanceRecommendation

Describes an Amazon EC2 instance recommendation.

Contents

accountId

The AWS account ID of the instance.

Type: String

Required: No

currentInstanceGpuInfo

Describes the GPU accelerator settings for the current instance type.

Type: [GpuInfo](#) object

Required: No

currentInstanceType

The instance type of the current instance.

Type: String

Required: No

currentPerformanceRisk

The risk of the current instance not meeting the performance needs of its workloads. The higher the risk, the more likely the current instance cannot meet the performance requirements of its workload.

Type: String

Valid Values: VeryLow | Low | Medium | High

Required: No

effectiveRecommendationPreferences

An object that describes the effective recommendation preferences for the instance.

Type: [EffectiveRecommendationPreferences](#) object

Required: No

externalMetricStatus

An object that describes Compute Optimizer's integration status with your external metrics provider.

Type: [ExternalMetricStatus](#) object

Required: No

finding

The finding classification of the instance.

Findings for instances include:

- **Underprovisioned** —An instance is considered under-provisioned when at least one specification of your instance, such as CPU, memory, or network, does not meet the performance requirements of your workload. Under-provisioned instances may lead to poor application performance.
- **Overprovisioned** —An instance is considered over-provisioned when at least one specification of your instance, such as CPU, memory, or network, can be sized down while still meeting the performance requirements of your workload, and no specification is under-provisioned. Over-provisioned instances may lead to unnecessary infrastructure cost.
- **Optimized** —An instance is considered optimized when all specifications of your instance, such as CPU, memory, and network, meet the performance requirements of your workload and is not over provisioned. For optimized resources, AWS Compute Optimizer might recommend a new generation instance type.

Note

The valid values in your API responses appear as OVER_PROVISIONED, UNDER_PROVISIONED, or OPTIMIZED.

Type: String

Valid Values: Underprovisioned | Overprovisioned | Optimized | NotOptimized

Required: No

findingReasonCodes

The reason for the finding classification of the instance.

Finding reason codes for instances include:

- **CPUOverprovisioned** — The instance's CPU configuration can be sized down while still meeting the performance requirements of your workload. This is identified by analyzing the CPUUtilization metric of the current instance during the look-back period.
- **CPUUnderprovisioned** — The instance's CPU configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better CPU performance. This is identified by analyzing the CPUUtilization metric of the current instance during the look-back period.
- **MemoryOverprovisioned** — The instance's memory configuration can be sized down while still meeting the performance requirements of your workload. This is identified by analyzing the memory utilization metric of the current instance during the look-back period.
- **MemoryUnderprovisioned** — The instance's memory configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better memory performance. This is identified by analyzing the memory utilization metric of the current instance during the look-back period.

Note

Memory utilization is analyzed only for resources that have the unified CloudWatch agent installed on them. For more information, see [Enabling memory utilization with the Amazon CloudWatch Agent](#) in the *AWS Compute Optimizer User Guide*. On Linux instances, Compute Optimizer analyses the mem_used_percent metric in the CWAgent namespace, or the legacy MemoryUtilization metric in the System/Linux namespace. On Windows instances, Compute Optimizer analyses the Memory % Committed Bytes In Use metric in the CWAgent namespace.

- **EBSThroughputOverprovisioned** — The instance's EBS throughput configuration can be sized down while still meeting the performance requirements of your workload. This is identified by analyzing the VolumeReadBytes and VolumeWriteBytes metrics of EBS volumes attached to the current instance during the look-back period.
- **EBSThroughputUnderprovisioned** — The instance's EBS throughput configuration doesn't meet the performance requirements of your workload and there is an alternative

instance type that provides better EBS throughput performance. This is identified by analyzing the `VolumeReadBytes` and `VolumeWriteBytes` metrics of EBS volumes attached to the current instance during the look-back period.

- **EBSIOPSOverprovisioned** — The instance's EBS IOPS configuration can be sized down while still meeting the performance requirements of your workload. This is identified by analyzing the `VolumeReadOps` and `VolumeWriteOps` metric of EBS volumes attached to the current instance during the look-back period.
- **EBSIOPSUnderprovisioned** — The instance's EBS IOPS configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better EBS IOPS performance. This is identified by analyzing the `VolumeReadOps` and `VolumeWriteOps` metric of EBS volumes attached to the current instance during the look-back period.
- **NetworkBandwidthOverprovisioned** — The instance's network bandwidth configuration can be sized down while still meeting the performance requirements of your workload. This is identified by analyzing the `NetworkIn` and `NetworkOut` metrics of the current instance during the look-back period.
- **NetworkBandwidthUnderprovisioned** — The instance's network bandwidth configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better network bandwidth performance. This is identified by analyzing the `NetworkIn` and `NetworkOut` metrics of the current instance during the look-back period. This finding reason happens when the `NetworkIn` or `NetworkOut` performance of an instance is impacted.
- **NetworkPPSOverprovisioned** — The instance's network PPS (packets per second) configuration can be sized down while still meeting the performance requirements of your workload. This is identified by analyzing the `NetworkPacketsIn` and `NetworkPacketsIn` metrics of the current instance during the look-back period.
- **NetworkPPSUnderprovisioned** — The instance's network PPS (packets per second) configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better network PPS performance. This is identified by analyzing the `NetworkPacketsIn` and `NetworkPacketsIn` metrics of the current instance during the look-back period.
- **DiskIOPSOverprovisioned** — The instance's disk IOPS configuration can be sized down while still meeting the performance requirements of your workload. This is identified by analyzing the `DiskReadOps` and `DiskWriteOps` metrics of the current instance during the look-back period.

- **DiskIOPSUnderprovisioned** — The instance's disk IOPS configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better disk IOPS performance. This is identified by analyzing the `DiskReadOps` and `DiskWriteOps` metrics of the current instance during the look-back period.
- **DiskThroughputOverprovisioned** — The instance's disk throughput configuration can be sized down while still meeting the performance requirements of your workload. This is identified by analyzing the `DiskReadBytes` and `DiskWriteBytes` metrics of the current instance during the look-back period.
- **DiskThroughputUnderprovisioned** — The instance's disk throughput configuration doesn't meet the performance requirements of your workload and there is an alternative instance type that provides better disk throughput performance. This is identified by analyzing the `DiskReadBytes` and `DiskWriteBytes` metrics of the current instance during the look-back period.

 Note

For more information about instance metrics, see [List the available CloudWatch metrics for your instances](#) in the *Amazon Elastic Compute Cloud User Guide*. For more information about EBS volume metrics, see [Amazon CloudWatch metrics for Amazon EBS](#) in the *Amazon Elastic Compute Cloud User Guide*.

Type: Array of strings

Valid Values: `CPUOverprovisioned` | `CPUUnderprovisioned` | `MemoryOverprovisioned` | `MemoryUnderprovisioned` | `EBSThroughputOverprovisioned` | `EBSThroughputUnderprovisioned` | `EBSIOPSOverprovisioned` | `EBSIOPSUnderprovisioned` | `NetworkBandwidthOverprovisioned` | `NetworkBandwidthUnderprovisioned` | `NetworkPPSOverprovisioned` | `NetworkPPSUnderprovisioned` | `DiskIOPSOverprovisioned` | `DiskIOPSUnderprovisioned` | `DiskThroughputOverprovisioned` | `DiskThroughputUnderprovisioned` | `GPUUnderprovisioned` | `GPUOverprovisioned` | `GPUMemoryUnderprovisioned` | `GPUMemoryOverprovisioned`

Required: No

idle

Describes if an Amazon EC2 instance is idle.

Type: String

Valid Values: `True` | `False`

Required: No

inferredWorkloadTypes

The applications that might be running on the instance as inferred by Compute Optimizer.

Compute Optimizer can infer if one of the following applications might be running on the instance:

- `AmazonEmr` - Infers that Amazon EMR might be running on the instance.
- `ApacheCassandra` - Infers that Apache Cassandra might be running on the instance.
- `ApacheHadoop` - Infers that Apache Hadoop might be running on the instance.
- `Memcached` - Infers that Memcached might be running on the instance.
- `NGINX` - Infers that NGINX might be running on the instance.
- `PostgreSql` - Infers that PostgreSQL might be running on the instance.
- `Redis` - Infers that Redis might be running on the instance.
- `Kafka` - Infers that Kafka might be running on the instance.
- `SQLServer` - Infers that SQLServer might be running on the instance.

Type: Array of strings

Valid Values: `AmazonEmr` | `ApacheCassandra` | `ApacheHadoop` | `Memcached` | `Nginx` | `PostgreSql` | `Redis` | `Kafka` | `SQLServer`

Required: No

instanceArn

The Amazon Resource Name (ARN) of the current instance.

Type: String

Required: No

instanceName

The name of the current instance.

Type: String

Required: No

instanceState

The state of the instance when the recommendation was generated.

Type: String

Valid Values: pending | running | shutting-down | terminated | stopping | stopped

Required: No

lastRefreshTimestamp

The timestamp of when the instance recommendation was last generated.

Type: Timestamp

Required: No

lookBackPeriodInDays

The number of days for which utilization metrics were analyzed for the instance.

Type: Double

Required: No

recommendationOptions

An array of objects that describe the recommendation options for the instance.

Type: Array of [InstanceRecommendationOption](#) objects

Required: No

recommendationSources

An array of objects that describe the source resource of the recommendation.

Type: Array of [RecommendationSource](#) objects

Required: No

tags

A list of tags assigned to your Amazon EC2 instance recommendations.

Type: Array of [Tag](#) objects

Required: No

utilizationMetrics

An array of objects that describe the utilization metrics of the instance.

Type: Array of [UtilizationMetric](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

InstanceRecommendationOption

Describes a recommendation option for an Amazon EC2 instance.

Contents

instanceGpuInfo

Describes the GPU accelerator settings for the recommended instance type.

Type: [GpuInfo](#) object

Required: No

instanceType

The instance type of the instance recommendation.

Type: String

Required: No

migrationEffort

The level of effort required to migrate from the current instance type to the recommended instance type.

For example, the migration effort is `Low` if Amazon EMR is the inferred workload type and an AWS Graviton instance type is recommended. The migration effort is `Medium` if a workload type couldn't be inferred but an AWS Graviton instance type is recommended. The migration effort is `VeryLow` if both the current and recommended instance types are of the same CPU architecture.

Type: String

Valid Values: `VeryLow` | `Low` | `Medium` | `High`

Required: No

performanceRisk

The performance risk of the instance recommendation option.

Performance risk indicates the likelihood of the recommended instance type not meeting the resource needs of your workload. Compute Optimizer calculates an individual performance risk score for each specification of the recommended instance, including CPU, memory, EBS throughput, EBS IOPS, disk throughput, disk IOPS, network throughput, and network PPS. The performance risk of the recommended instance is calculated as the maximum performance risk score across the analyzed resource specifications.

The value ranges from 0 - 4, with 0 meaning that the recommended resource is predicted to always provide enough hardware capability. The higher the performance risk is, the more likely you should validate whether the recommendation will meet the performance requirements of your workload before migrating your resource.

Type: Double

Valid Range: Minimum value of 0. Maximum value of 4.

Required: No

platformDifferences

Describes the configuration differences between the current instance and the recommended instance type. You should consider the configuration differences before migrating your workloads from the current instance to the recommended instance type. The [Change the instance type guide for Linux](#) and [Change the instance type guide for Windows](#) provide general guidance for getting started with an instance migration.

Platform differences include:

- **Hypervisor** — The hypervisor of the recommended instance type is different than that of the current instance. For example, the recommended instance type uses a Nitro hypervisor and the current instance uses a Xen hypervisor. The differences that you should consider between these hypervisors are covered in the [Nitro Hypervisor](#) section of the Amazon EC2 frequently asked questions. For more information, see [Instances built on the Nitro System](#) in the *Amazon EC2 User Guide for Linux*, or [Instances built on the Nitro System](#) in the *Amazon EC2 User Guide for Windows*.
- **NetworkInterface** — The network interface of the recommended instance type is different than that of the current instance. For example, the recommended instance type supports enhanced networking and the current instance might not. To enable enhanced networking for the recommended instance type, you must install the Elastic Network Adapter (ENA) driver or the Intel 82599 Virtual Function driver. For more information, see [Networking](#)

[and storage features](#) and [Enhanced networking on Linux](#) in the *Amazon EC2 User Guide for Linux*, or [Networking and storage features](#) and [Enhanced networking on Windows](#) in the *Amazon EC2 User Guide for Windows*.

- **StorageInterface** — The storage interface of the recommended instance type is different than that of the current instance. For example, the recommended instance type uses an NVMe storage interface and the current instance does not. To access NVMe volumes for the recommended instance type, you will need to install or upgrade the NVMe driver. For more information, see [Networking and storage features](#) and [Amazon EBS and NVMe on Linux instances](#) in the *Amazon EC2 User Guide for Linux*, or [Networking and storage features](#) and [Amazon EBS and NVMe on Windows instances](#) in the *Amazon EC2 User Guide for Windows*.
- **InstanceStoreAvailability** — The recommended instance type does not support instance store volumes and the current instance does. Before migrating, you might need to back up the data on your instance store volumes if you want to preserve them. For more information, see [How do I back up an instance store volume on my Amazon EC2 instance to Amazon EBS?](#) in the *AWS Premium Support Knowledge Base*. For more information, see [Networking and storage features](#) and [Amazon EC2 instance store](#) in the *Amazon EC2 User Guide for Linux*, or see [Networking and storage features](#) and [Amazon EC2 instance store](#) in the *Amazon EC2 User Guide for Windows*.
- **VirtualizationType** — The recommended instance type uses the hardware virtual machine (HVM) virtualization type and the current instance uses the paravirtual (PV) virtualization type. For more information about the differences between these virtualization types, see [Linux AMI virtualization types](#) in the *Amazon EC2 User Guide for Linux*, or [Windows AMI virtualization types](#) in the *Amazon EC2 User Guide for Windows*.
- **Architecture** — The CPU architecture between the recommended instance type and the current instance is different. For example, the recommended instance type might use an Arm CPU architecture and the current instance type might use a different one, such as x86. Before migrating, you should consider recompiling the software on your instance for the new architecture. Alternatively, you might switch to an Amazon Machine Image (AMI) that supports the new architecture. For more information about the CPU architecture for each instance type, see [Amazon EC2 Instance Types](#).

Type: Array of strings

Valid Values: Hypervisor | NetworkInterface | StorageInterface | InstanceStoreAvailability | VirtualizationType | Architecture

Required: No

projectedUtilizationMetrics

An array of objects that describe the projected utilization metrics of the instance recommendation option.

Note

The Cpu and Memory metrics are the only projected utilization metrics returned. Additionally, the Memory metric is returned only for resources that have the unified CloudWatch agent installed on them. For more information, see [Enabling Memory Utilization with the CloudWatch Agent](#).

Type: Array of [UtilizationMetric](#) objects

Required: No

rank

The rank of the instance recommendation option.

The top recommendation option is ranked as 1.

Type: Integer

Required: No

savingsOpportunity

An object that describes the savings opportunity for the instance recommendation option. Savings opportunity includes the estimated monthly savings amount and percentage.

Type: [SavingsOpportunity](#) object

Required: No

savingsOpportunityAfterDiscounts

An object that describes the savings opportunity for the instance recommendation option that includes Savings Plans and Reserved Instances discounts. Savings opportunity includes the estimated monthly savings and percentage.

Type: [InstanceSavingsOpportunityAfterDiscounts](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

InstanceSavingsEstimationMode

Describes the savings estimation mode used for calculating savings opportunity for Amazon EC2 instances.

Contents

source

Describes the source for calculating the savings opportunity for Amazon EC2 instances.

Type: String

Valid Values: `PublicPricing` | `CostExplorerRightsizing` | `CostOptimizationHub`

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

InstanceSavingsOpportunityAfterDiscounts

Describes the savings opportunity for instance recommendations after applying the Savings Plans and Reserved Instances discounts.

Savings opportunity after discounts represents the estimated monthly savings you can achieve by implementing Compute Optimizer recommendations.

Contents

estimatedMonthlySavings

An object that describes the estimated monthly savings possible by adopting Compute Optimizer's Amazon EC2 instance recommendations. This is based on pricing after applying the Savings Plans and Reserved Instances discounts.

Type: [InstanceEstimatedMonthlySavings](#) object

Required: No

savingsOpportunityPercentage

The estimated monthly savings possible as a percentage of monthly cost after applying the Savings Plans and Reserved Instances discounts. This saving can be achieved by adopting Compute Optimizer's EC2 instance recommendations.

Type: Double

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

JobFilter

Describes a filter that returns a more specific list of recommendation export jobs. Use this filter with the [DescribeRecommendationExportJobs](#) action.

You can use `EBSFilter` with the [GetEBSVolumeRecommendations](#) action, `LambdaFunctionRecommendationFilter` with the [GetLambdaFunctionRecommendations](#) action, and `Filter` with the [GetAutoScalingGroupRecommendations](#) and [GetEC2InstanceRecommendations](#) actions.

Contents

name

The name of the filter.

Specify `ResourceType` to return export jobs of a specific resource type (for example, `Ec2Instance`).

Specify `JobStatus` to return export jobs with a specific status (e.g, `Complete`).

Type: String

Valid Values: `ResourceType` | `JobStatus`

Required: No

values

The value of the filter.

The valid values for this parameter are as follows, depending on what you specify for the `name` parameter:

- Specify `Ec2Instance` or `AutoScalingGroup` if you specify the `name` parameter as `ResourceType`. There is no filter for EBS volumes because volume recommendations cannot be exported at this time.
- Specify `Queued`, `InProgress`, `Complete`, or `Failed` if you specify the `name` parameter as `JobStatus`.

Type: Array of strings

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

LambdaEffectiveRecommendationPreferences

Describes the effective recommendation preferences for Lambda functions.

Contents

savingsEstimationMode

Describes the savings estimation mode applied for calculating savings opportunity for Lambda functions.

Type: [LambdaSavingsEstimationMode](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

LambdaEstimatedMonthlySavings

Describes the estimated monthly savings possible for Lambda functions by adopting Compute Optimizer recommendations. This is based on Lambda functions pricing after applying Savings Plans discounts.

Contents

currency

The currency of the estimated monthly savings.

Type: String

Valid Values: USD | CNY

Required: No

value

The value of the estimated monthly savings.

Type: Double

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

LambdaFunctionMemoryProjectedMetric

Describes a projected utilization metric of an AWS Lambda function recommendation option.

Contents

name

The name of the projected utilization metric.

Type: String

Valid Values: Duration

Required: No

statistic

The statistic of the projected utilization metric.

Type: String

Valid Values: LowerBound | UpperBound | Expected

Required: No

value

The values of the projected utilization metrics.

Type: Double

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

LambdaFunctionMemoryRecommendationOption

Describes a recommendation option for an AWS Lambda function.

Contents

memorySize

The memory size, in MB, of the function recommendation option.

Type: Integer

Required: No

projectedUtilizationMetrics

An array of objects that describe the projected utilization metrics of the function recommendation option.

Type: Array of [LambdaFunctionMemoryProjectedMetric](#) objects

Required: No

rank

The rank of the function recommendation option.

The top recommendation option is ranked as 1.

Type: Integer

Required: No

savingsOpportunity

An object that describes the savings opportunity for the Lambda function recommendation option. Savings opportunity includes the estimated monthly savings amount and percentage.

Type: [SavingsOpportunity](#) object

Required: No

savingsOpportunityAfterDiscounts

An object that describes the savings opportunity for the Lambda recommendation option which includes Saving Plans discounts. Savings opportunity includes the estimated monthly savings and percentage.

Type: [LambdaSavingsOpportunityAfterDiscounts](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

LambdaFunctionRecommendation

Describes an AWS Lambda function recommendation.

Contents

accountId

The AWS account ID of the function.

Type: String

Required: No

currentMemorySize

The amount of memory, in MB, that's allocated to the current function.

Type: Integer

Required: No

currentPerformanceRisk

The risk of the current Lambda function not meeting the performance needs of its workloads. The higher the risk, the more likely the current Lambda function requires more memory.

Type: String

Valid Values: VeryLow | Low | Medium | High

Required: No

effectiveRecommendationPreferences

Describes the effective recommendation preferences for Lambda functions.

Type: [LambdaEffectiveRecommendationPreferences](#) object

Required: No

finding

The finding classification of the function.

Findings for functions include:

- **Optimized** — The function is correctly provisioned to run your workload based on its current configuration and its utilization history. This finding classification does not include finding reason codes.
- **NotOptimized** — The function is performing at a higher level (over-provisioned) or at a lower level (under-provisioned) than required for your workload because its current configuration is not optimal. Over-provisioned resources might lead to unnecessary infrastructure cost, and under-provisioned resources might lead to poor application performance. This finding classification can include the `MemoryUnderprovisioned` and `MemoryUnderprovisioned` finding reason codes.
- **Unavailable** — Compute Optimizer was unable to generate a recommendation for the function. This could be because the function has not accumulated sufficient metric data, or the function does not qualify for a recommendation. This finding classification can include the `InsufficientData` and `Inconclusive` finding reason codes.

 Note

Functions with a finding of unavailable are not returned unless you specify the `filter` parameter with a value of `Unavailable` in your `GetLambdaFunctionRecommendations` request.

Type: String

Valid Values: `Optimized` | `NotOptimized` | `Unavailable`

Required: No

findingReasonCodes

The reason for the finding classification of the function.

 Note

Functions that have a finding classification of `Optimized` don't have a finding reason code.

Finding reason codes for functions include:

- **MemoryOverprovisioned** — The function is over-provisioned when its memory configuration can be sized down while still meeting the performance requirements of your workload. An over-provisioned function might lead to unnecessary infrastructure cost. This finding reason code is part of the `NotOptimized` finding classification.
- **MemoryUnderprovisioned** — The function is under-provisioned when its memory configuration doesn't meet the performance requirements of the workload. An under-provisioned function might lead to poor application performance. This finding reason code is part of the `NotOptimized` finding classification.
- **InsufficientData** — The function does not have sufficient metric data for Compute Optimizer to generate a recommendation. For more information, see the [Supported resources and requirements](#) in the *AWS Compute Optimizer User Guide*. This finding reason code is part of the `Unavailable` finding classification.
- **Inconclusive** — The function does not qualify for a recommendation because Compute Optimizer cannot generate a recommendation with a high degree of confidence. This finding reason code is part of the `Unavailable` finding classification.

Type: Array of strings

Valid Values: `MemoryOverprovisioned` | `MemoryUnderprovisioned` | `InsufficientData` | `Inconclusive`

Required: No

functionArn

The Amazon Resource Name (ARN) of the current function.

Type: String

Required: No

functionVersion

The version number of the current function.

Type: String

Required: No

lastRefreshTimestamp

The timestamp of when the function recommendation was last generated.

Type: Timestamp

Required: No

lookbackPeriodInDays

The number of days for which utilization metrics were analyzed for the function.

Type: Double

Required: No

memorySizeRecommendationOptions

An array of objects that describe the memory configuration recommendation options for the function.

Type: Array of [LambdaFunctionMemoryRecommendationOption](#) objects

Required: No

numberOfInvocations

The number of times your function code was applied during the look-back period.

Type: Long

Required: No

tags

A list of tags assigned to your Lambda function recommendations.

Type: Array of [Tag](#) objects

Required: No

utilizationMetrics

An array of objects that describe the utilization metrics of the function.

Type: Array of [LambdaFunctionUtilizationMetric](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

LambdaFunctionRecommendationFilter

Describes a filter that returns a more specific list of AWS Lambda function recommendations. Use this filter with the [GetLambdaFunctionRecommendations](#) action.

You can use `EBSFilter` with the [GetEBSVolumeRecommendations](#) action, `JobFilter` with the [DescribeRecommendationExportJobs](#) action, and `Filter` with the [GetAutoScalingGroupRecommendations](#) and [GetEC2InstanceRecommendations](#) actions.

Contents

name

The name of the filter.

Specify `Finding` to return recommendations with a specific finding classification (for example, `NotOptimized`).

Specify `FindingReasonCode` to return recommendations with a specific finding reason code (for example, `MemoryUnderprovisioned`).

You can filter your Lambda function recommendations by `tag:key` and `tag-key` tags.

A `tag:key` is a key and value combination of a tag assigned to your Lambda function recommendations. Use the tag key in the filter name and the tag value as the filter value. For example, to find all Lambda function recommendations that have a tag with the key of `Owner` and the value of `TeamA`, specify `tag:Owner` for the filter name and `TeamA` for the filter value.

A `tag-key` is the key of a tag assigned to your Lambda function recommendations. Use this filter to find all of your Lambda function recommendations that have a tag with a specific key. This doesn't consider the tag value. For example, you can find your Lambda function recommendations with a tag key value of `Owner` or without any tag keys assigned.

Type: String

Valid Values: `Finding` | `FindingReasonCode`

Required: No

values

The value of the filter.

The valid values for this parameter are as follows, depending on what you specify for the name parameter:

- Specify `Optimized`, `NotOptimized`, or `Unavailable` if you specify the name parameter as `Finding`.
- Specify `MemoryOverprovisioned`, `MemoryUnderprovisioned`, `InsufficientData`, or `Inconclusive` if you specify the name parameter as `FindingReasonCode`.

Type: Array of strings

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

LambdaFunctionUtilizationMetric

Describes a utilization metric of an AWS Lambda function.

Contents

name

The name of the utilization metric.

The following utilization metrics are available:

- `Duration` - The amount of time that your function code spends processing an event.
- `Memory` - The amount of memory used per invocation.

Type: String

Valid Values: `Duration` | `Memory`

Required: No

statistic

The statistic of the utilization metric.

The Compute Optimizer API, AWS Command Line Interface (AWS CLI), and SDKs return utilization metrics using only the `Maximum` statistic, which is the highest value observed during the specified period.

The Compute Optimizer console displays graphs for some utilization metrics using the `Average` statistic, which is the value of `Sum / SampleCount` during the specified period. For more information, see [Viewing resource recommendations](#) in the *AWS Compute Optimizer User Guide*. You can also get averaged utilization metric data for your resources using Amazon CloudWatch. For more information, see the [Amazon CloudWatch User Guide](#).

Type: String

Valid Values: `Maximum` | `Average`

Required: No

value

The value of the utilization metric.

Type: Double

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

LambdaSavingsEstimationMode

Describes the savings estimation used for calculating savings opportunity for Lambda functions.

Contents

source

Describes the source for calculation of savings opportunity for Lambda functions.

Type: String

Valid Values: `PublicPricing` | `CostExplorerRightsizing` | `CostOptimizationHub`

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

LambdaSavingsOpportunityAfterDiscounts

Describes the savings opportunity for Lambda functions recommendations after applying Savings Plans discounts.

Savings opportunity represents the estimated monthly savings after applying Savings Plans discounts. You can achieve this by implementing a given Compute Optimizer recommendation.

Contents

estimatedMonthlySavings

The estimated monthly savings possible by adopting Compute Optimizer's Lambda function recommendations. This includes any applicable Savings Plans discounts.

Type: [LambdaEstimatedMonthlySavings](#) object

Required: No

savingsOpportunityPercentage

The estimated monthly savings possible as a percentage of monthly cost by adopting Compute Optimizer's Lambda function recommendations. This includes any applicable Savings Plans discounts.

Type: Double

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

LicenseConfiguration

Describes the configuration of a license for an Amazon EC2 instance.

Contents

instanceType

The instance type used in the license.

Type: String

Required: No

licenseEdition

The edition of the license for the application that runs on the instance.

Type: String

Valid Values: Enterprise | Standard | Free | NoLicenseEditionFound

Required: No

licenseModel

The license type associated with the instance.

Type: String

Valid Values: LicenseIncluded | BringYourOwnLicense

Required: No

licenseName

The name of the license for the application that runs on the instance.

Type: String

Valid Values: SQLServer

Required: No

licenseVersion

The version of the license for the application that runs on the instance.

Type: String

Required: No

metricsSource

The list of metric sources required to generate recommendations for commercial software licenses.

Type: Array of [MetricSource](#) objects

Required: No

numberOfCores

The current number of cores associated with the instance.

Type: Integer

Required: No

operatingSystem

The operating system of the instance.

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

LicenseRecommendation

Describes a license recommendation for an EC2 instance.

Contents

accountId

The AWS account ID of the license.

Type: String

Required: No

currentLicenseConfiguration

An object that describes the current configuration of an instance that runs on a license.

Type: [LicenseConfiguration](#) object

Required: No

finding

The finding classification for an instance that runs on a license.

Findings include:

- **InsufficientMetrics** — When Compute Optimizer detects that your CloudWatch Application Insights isn't enabled or is enabled with insufficient permissions.
- **NotOptimized** — When Compute Optimizer detects that your EC2 infrastructure isn't using any of the SQL server license features you're paying for, a license is considered not optimized.
- **Optimized** — When Compute Optimizer detects that all specifications of your license meet the performance requirements of your workload.

Type: String

Valid Values: `InsufficientMetrics` | `Optimized` | `NotOptimized`

Required: No

findingReasonCodes

The reason for the finding classification for an instance that runs on a license.

Finding reason codes include:

- **Optimized** — All specifications of your license meet the performance requirements of your workload.
- **LicenseOverprovisioned** — A license is considered over-provisioned when your license can be downgraded while still meeting the performance requirements of your workload.
- **InvalidCloudwatchApplicationInsights** — CloudWatch Application Insights isn't configured properly.
- **CloudwatchApplicationInsightsError** — There is a CloudWatch Application Insights error.

Type: Array of strings

Valid Values: `InvalidCloudWatchApplicationInsightsSetup` | `CloudWatchApplicationInsightsError` | `LicenseOverprovisioned` | `Optimized`

Required: No

lastRefreshTimestamp

The timestamp of when the license recommendation was last generated.

Type: Timestamp

Required: No

licenseRecommendationOptions

An array of objects that describe the license recommendation options.

Type: Array of [LicenseRecommendationOption](#) objects

Required: No

lookbackPeriodInDays

The number of days for which utilization metrics were analyzed for an instance that runs on a license.

Type: Double

Required: No

resourceArn

The ARN that identifies the Amazon EC2 instance.

Type: String

Required: No

tags

A list of tags assigned to an EC2 instance.

Type: Array of [Tag](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

LicenseRecommendationFilter

Describes a filter that returns a more specific list of license recommendations. Use this filter with the `GetLicenseRecommendation` action.

Contents

name

The name of the filter.

Specify `Finding` to return recommendations with a specific finding classification.

Specify `FindingReasonCode` to return recommendations with a specific finding reason code.

You can filter your license recommendations by `tag:key` and `tag-key` tags.

A `tag:key` is a key and value combination of a tag assigned to your license recommendations. Use the `tag key` in the filter name and the tag value as the filter value. For example, to find all license recommendations that have a tag with the key of `Owner` and the value of `TeamA`, specify `tag:Owner` for the filter name and `TeamA` for the filter value.

A `tag-key` is the key of a tag assigned to your license recommendations. Use this filter to find all of your license recommendations that have a tag with a specific key. This doesn't consider the tag value. For example, you can find your license recommendations with a tag key value of `Owner` or without any tag keys assigned.

Type: String

Valid Values: `Finding` | `FindingReasonCode` | `LicenseName`

Required: No

values

The value of the filter.

The valid values for this parameter are as follows, depending on what you specify for the `name` parameter:

- If you specify the `name` parameter as `Finding`, then specify `Optimized`, `NotOptimized`, or `InsufficientMetrics`.

- If you specify the name parameter as `FindingReasonCode`, then specify `Optimized`, `LicenseOverprovisioned`, `InvalidCloudwatchApplicationInsights`, or `CloudwatchApplicationInsightsError`.

Type: Array of strings

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

LicenseRecommendationOption

Describes the recommendation options for licenses.

Contents

licenseEdition

The recommended edition of the license for the application that runs on the instance.

Type: String

Valid Values: Enterprise | Standard | Free | NoLicenseEditionFound

Required: No

licenseModel

The recommended license type associated with the instance.

Type: String

Valid Values: LicenseIncluded | BringYourOwnLicense

Required: No

operatingSystem

The operating system of a license recommendation option.

Type: String

Required: No

rank

The rank of the license recommendation option.

The top recommendation option is ranked as 1.

Type: Integer

Required: No

savingsOpportunity

Describes the savings opportunity for recommendations of a given resource type or for the recommendation option of an individual resource.

Savings opportunity represents the estimated monthly savings you can achieve by implementing a given Compute Optimizer recommendation.

Important

Savings opportunity data requires that you opt in to Cost Explorer, as well as activate **Receive Amazon EC2 resource recommendations** in the Cost Explorer preferences page. That creates a connection between Cost Explorer and Compute Optimizer. With this connection, Cost Explorer generates savings estimates considering the price of existing resources, the price of recommended resources, and historical usage data. Estimated monthly savings reflects the projected dollar savings associated with each of the recommendations generated. For more information, see [Enabling Cost Explorer](#) and [Optimizing your cost with Rightsizing Recommendations](#) in the *Cost Management User Guide*.

Type: [SavingsOpportunity](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

MemorySizeConfiguration

The memory size configurations of a container.

Contents

memory

The amount of memory in the container.

Type: Integer

Required: No

memoryReservation

The limit of memory reserve for the container.

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

MetricSource

The list of metric sources required to generate recommendations for commercial software licenses.

Contents

provider

The name of the metric source provider.

Type: String

Valid Values: CloudWatchApplicationInsights

Required: No

providerArn

The ARN of the metric source provider.

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

OrderBy

Describes how the recommendations are ordered.

Contents

dimension

The dimension values to sort the recommendations.

Type: String

Valid Values: SavingsValue | SavingsValueAfterDiscount

Required: No

order

The order to sort the recommendations.

Type: String

Valid Values: Asc | Desc

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

PreferredResource

The preference to control which resource type values are considered when generating rightsizing recommendations. You can specify this preference as a combination of include and exclude lists. You must specify either an `includeList` or `excludeList`. If the preference is an empty set of resource type values, an error occurs. For more information, see [Rightsizing recommendation preferences](#) in the *AWS Compute Optimizer User Guide*.

Note

- This preference is only available for the Amazon EC2 instance and Auto Scaling group resource types.
- Compute Optimizer only supports the customization of `Ec2InstanceTypes`.

Contents

excludeList

The preferred resource type values to exclude from the recommendation candidates. If this isn't specified, all supported resources are included by default. You can specify up to 1000 values in this list.

Type: Array of strings

Required: No

includeList

The preferred resource type values to include in the recommendation candidates. You can specify the exact resource type value, such as `m5.large`, or use wild card expressions, such as `m5`. If this isn't specified, all supported resources are included by default. You can specify up to 1000 values in this list.

Type: Array of strings

Required: No

name

The type of preferred resource to customize.

Note

Compute Optimizer only supports the customization of `Ec2InstanceTypes`.

Type: String

Valid Values: `Ec2InstanceTypes`

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

ProjectedMetric

Describes a projected utilization metric of a recommendation option, such as an Amazon EC2 instance. This represents the projected utilization of a recommendation option had you used that resource during the analyzed period.

Compare the utilization metric data of your resource against its projected utilization metric data to determine the performance difference between your current resource and the recommended option.

Note

The Cpu, Memory, GPU, and GPU_MEMORY metrics are the only projected utilization metrics returned when you run the [GetEC2RecommendationProjectedMetrics](#) action. Additionally, these metrics are only returned for resources with the unified CloudWatch agent installed on them. For more information, see [Enabling Memory Utilization with the CloudWatch Agent](#) and [Enabling NVIDIA GPU utilization with the CloudWatch Agent](#).

Contents

name

The name of the projected utilization metric.

The following projected utilization metrics are returned:

- Cpu - The projected percentage of allocated EC2 compute units that would be in use on the recommendation option had you used that resource during the analyzed period. This metric identifies the processing power required to run an application on the recommendation option.

Depending on the instance type, tools in your operating system can show a lower percentage than CloudWatch when the instance is not allocated a full processor core.

- Memory - The percentage of memory that would be in use on the recommendation option had you used that resource during the analyzed period. This metric identifies the amount of memory required to run an application on the recommendation option.

Units: Percent

Note

The Memory metric is only returned for resources with the unified CloudWatch agent installed on them. For more information, see [Enabling Memory Utilization with the CloudWatch Agent](#).

- GPU - The projected percentage of allocated GPUs if you adjust your configurations to Compute Optimizer's recommendation option.
- GPU_MEMORY - The projected percentage of total GPU memory if you adjust your configurations to Compute Optimizer's recommendation option.

Note

The GPU and GPU_MEMORY metrics are only returned for resources with the unified CloudWatch Agent installed on them. For more information, see [Enabling NVIDIA GPU utilization with the CloudWatch Agent](#).

Type: String

Valid Values: Cpu | Memory | EBS_READ_OPS_PER_SECOND | EBS_WRITE_OPS_PER_SECOND | EBS_READ_BYTES_PER_SECOND | EBS_WRITE_BYTES_PER_SECOND | DISK_READ_OPS_PER_SECOND | DISK_WRITE_OPS_PER_SECOND | DISK_READ_BYTES_PER_SECOND | DISK_WRITE_BYTES_PER_SECOND | NETWORK_IN_BYTES_PER_SECOND | NETWORK_OUT_BYTES_PER_SECOND | NETWORK_PACKETS_IN_PER_SECOND | NETWORK_PACKETS_OUT_PER_SECOND | GPU_PERCENTAGE | GPU_MEMORY_PERCENTAGE

Required: No

timestamps

The timestamps of the projected utilization metric.

Type: Array of timestamps

Required: No

values

The values of the projected utilization metrics.

Type: Array of doubles

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

RDSDatabaseProjectedMetric

Describes the projected metrics of an Amazon Aurora and RDS database recommendation option.

To determine the performance difference between your current Amazon Aurora and RDS database and the recommended option, compare the metric data of your service against its projected metric data.

Contents

name

The name of the projected metric.

Type: String

Valid Values: CPU | Memory | EBSVolumeStorageSpaceUtilization | NetworkReceiveThroughput | NetworkTransmitThroughput | EBSVolumeReadIOPS | EBSVolumeWriteIOPS | EBSVolumeReadThroughput | EBSVolumeWriteThroughput | DatabaseConnections | StorageNetworkReceiveThroughput | StorageNetworkTransmitThroughput | AuroraMemoryHealthState | AuroraMemoryNumDeclinedSql | AuroraMemoryNumKillConnTotal | AuroraMemoryNumKillQueryTotal | ReadIOPSEphemeralStorage | WriteIOPSEphemeralStorage | VolumeReadIOPs | VolumeBytesUsed | VolumeWriteIOPs

Required: No

timestamps

The timestamps of the projected metric.

Type: Array of timestamps

Required: No

values

The values for the projected metric.

Type: Array of doubles

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

RDSDatabaseRecommendedOptionProjectedMetric

Describes the projected metrics of an Amazon Aurora and RDS database recommendation option.

To determine the performance difference between your current Amazon Aurora and RDS database and the recommended option, compare the metric data of your service against its projected metric data.

Contents

projectedMetrics

An array of objects that describe the projected metric.

Type: Array of [RDSDatabaseProjectedMetric](#) objects

Required: No

rank

The rank identifier of the Amazon Aurora or RDS DB instance recommendation option.

Type: Integer

Required: No

recommendedDBInstanceClass

The recommended DB instance class for the Amazon Aurora or RDS database.

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

RDSDBInstanceRecommendationOption

Describes the recommendation options for a DB instance.

Contents

dbInstanceClass

Describes the DB instance class recommendation option for your Amazon Aurora or RDS database.

Type: String

Required: No

performanceRisk

The performance risk of the DB instance recommendation option.

Type: Double

Valid Range: Minimum value of 0. Maximum value of 4.

Required: No

projectedUtilizationMetrics

An array of objects that describe the projected utilization metrics of the DB instance recommendation option.

Type: Array of [RDSDBUtilizationMetric](#) objects

Required: No

rank

The rank identifier of the DB instance recommendation option.

Type: Integer

Required: No

savingsOpportunity

Describes the savings opportunity for recommendations of a given resource type or for the recommendation option of an individual resource.

Savings opportunity represents the estimated monthly savings you can achieve by implementing a given Compute Optimizer recommendation.

⚠ Important

Savings opportunity data requires that you opt in to Cost Explorer, as well as activate **Receive Amazon EC2 resource recommendations** in the Cost Explorer preferences page. That creates a connection between Cost Explorer and Compute Optimizer. With this connection, Cost Explorer generates savings estimates considering the price of existing resources, the price of recommended resources, and historical usage data. Estimated monthly savings reflects the projected dollar savings associated with each of the recommendations generated. For more information, see [Enabling Cost Explorer](#) and [Optimizing your cost with Rightsizing Recommendations](#) in the *Cost Management User Guide*.

Type: [SavingsOpportunity](#) object

Required: No

savingsOpportunityAfterDiscounts

Describes the savings opportunity for Amazon Aurora and RDS database recommendations or for the recommendation option.

Savings opportunity represents the estimated monthly savings after applying Savings Plans discounts. You can achieve this by implementing a given Compute Optimizer recommendation.

Type: [RDSInstanceSavingsOpportunityAfterDiscounts](#) 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 Java V2](#)

- [AWS SDK for Ruby V3](#)

RDSDBRecommendation

Describes an Amazon Aurora and RDS database recommendation.

Contents

accountId

The AWS account ID of the Amazon Aurora or RDS database.

Type: String

Required: No

currentDBInstanceClass

The DB instance class of the current Aurora or RDS DB instance.

Type: String

Required: No

currentInstancePerformanceRisk

The performance risk for the current DB instance.

Type: String

Valid Values: VeryLow | Low | Medium | High

Required: No

currentStorageConfiguration

The configuration of the current DB storage.

Type: [DBStorageConfiguration](#) object

Required: No

currentStorageEstimatedMonthlyVolumeIOPsCostVariation

The level of variation in monthly I/O costs for the current DB storage configuration.

Type: String

Valid Values: None | Low | Medium | High

Required: No

dbClusterIdentifier

The identifier for DB cluster.

Type: String

Required: No

effectiveRecommendationPreferences

Describes the effective recommendation preferences for DB instances.

Type: [RDSEffectiveRecommendationPreferences](#) object

Required: No

engine

The engine of the DB instance.

Type: String

Required: No

engineVersion

The database engine version.

Type: String

Required: No

idle

This indicates if the DB instance is idle or not.

Type: String

Valid Values: True | False

Required: No

instanceFinding

The finding classification of an Amazon Aurora and RDS DB instance.

For more information about finding classifications, see [Finding classifications for Aurora and RDS databases](#) in the *AWS Compute Optimizer User Guide*.

Type: String

Valid Values: Optimized | Underprovisioned | Overprovisioned

Required: No

instanceFindingReasonCodes

The reason for the finding classification of a DB instance.

Type: Array of strings

Valid Values: CPUOverprovisioned | NetworkBandwidthOverprovisioned
| EBSIOPSOverprovisioned | EBSIOPSUnderprovisioned |
EBSThroughputOverprovisioned | CPUUnderprovisioned |
NetworkBandwidthUnderprovisioned | EBSThroughputUnderprovisioned |
NewGenerationDBInstanceClassAvailable | NewEngineVersionAvailable
| DBClusterWriterUnderprovisioned | MemoryUnderprovisioned |
InstanceStorageReadIOPSUnderprovisioned |
InstanceStorageWriteIOPSUnderprovisioned

Required: No

instanceRecommendationOptions

An array of objects that describe the recommendation options for the RDS DB instance.

Type: Array of [RDSDBInstanceRecommendationOption](#) objects

Required: No

lastRefreshTimestamp

The timestamp of when the DB instance recommendation was last generated.

Type: Timestamp

Required: No

lookbackPeriodInDays

The number of days the DB instance utilization metrics were analyzed.

Type: Double

Required: No

promotionTier

The promotion tier for the Aurora instance.

Type: Integer

Required: No

resourceArn

The ARN of the current Amazon Aurora or RDS database.

The following is the format of the ARN:

```
arn:aws:rds:{region}:{accountId}:db:{resourceName}
```

Type: String

Required: No

storageFinding

The finding classification of Amazon RDS DB instance storage.

For more information about finding classifications, see [Finding classifications for Aurora and RDS databases](#) in the *AWS Compute Optimizer User Guide*.

Type: String

Valid Values: Optimized | Underprovisioned | Overprovisioned | NotOptimized

Required: No

storageFindingReasonCodes

The reason for the finding classification of RDS DB instance storage.

Type: Array of strings

Valid Values: EBSVolumeAllocatedStorageUnderprovisioned |
EBSVolumeThroughputUnderprovisioned | EBSVolumeIOPSOverprovisioned |

EBSVolumeThroughputOverprovisioned | NewGenerationStorageTypeAvailable |
DBClusterStorageOptionAvailable | DBClusterStorageSavingsAvailable

Required: No

storageRecommendationOptions

An array of objects that describe the recommendation options for DB instance storage.

Type: Array of [RDSDBStorageRecommendationOption](#) objects

Required: No

tags

A list of tags assigned to your DB instance recommendations.

Type: Array of [Tag](#) objects

Required: No

utilizationMetrics

An array of objects that describe the utilization metrics of the DB instance.

Type: Array of [RDSDBUtilizationMetric](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

RDSDBRecommendationFilter

Describes a filter that returns a more specific list of DB instance recommendations. Use this filter with the [GetECSServiceRecommendations](#) action.

Contents

name

The name of the filter.

Specify `Finding` to return recommendations with a specific finding classification.

You can filter your DB instance recommendations by `tag:key` and `tag-key` tags.

A `tag:key` is a key and value combination of a tag assigned to your DB instance recommendations. Use the tag key in the filter name and the tag value as the filter value. For example, to find all DB instance recommendations that have a tag with the key of `Owner` and the value of `TeamA`, specify `tag:Owner` for the filter name and `TeamA` for the filter value.

A `tag-key` is the key of a tag assigned to your DB instance recommendations. Use this filter to find all of your DB instance recommendations that have a tag with a specific key. This doesn't consider the tag value. For example, you can find your DB instance recommendations with a tag key value of `Owner` or without any tag keys assigned.

Type: String

Valid Values: `InstanceFinding` | `InstanceFindingReasonCode` | `StorageFinding` | `StorageFindingReasonCode` | `Idle`

Required: No

values

The value of the filter.

Type: Array of strings

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

RDSDBStorageRecommendationOption

Describes the recommendation options for DB storage.

Contents

estimatedMonthlyVolumeIOPsCostVariation

The projected level of variation in monthly I/O costs for the DB storage recommendation option.

Type: String

Valid Values: None | Low | Medium | High

Required: No

rank

The rank identifier of the DB storage recommendation option.

Type: Integer

Required: No

savingsOpportunity

Describes the savings opportunity for recommendations of a given resource type or for the recommendation option of an individual resource.

Savings opportunity represents the estimated monthly savings you can achieve by implementing a given Compute Optimizer recommendation.

Important

Savings opportunity data requires that you opt in to Cost Explorer, as well as activate **Receive Amazon EC2 resource recommendations** in the Cost Explorer preferences page. That creates a connection between Cost Explorer and Compute Optimizer. With this connection, Cost Explorer generates savings estimates considering the price of existing resources, the price of recommended resources, and historical usage data. Estimated monthly savings reflects the projected dollar savings associated with each of the recommendations generated. For more information, see [Enabling Cost Explorer](#) and

[Optimizing your cost with Rightsizing Recommendations](#) in the *Cost Management User Guide*.

Type: [SavingsOpportunity](#) object

Required: No

savingsOpportunityAfterDiscounts

Describes the savings opportunity for DB storage recommendations or for the recommendation option.

Savings opportunity represents the estimated monthly savings after applying Savings Plans discounts. You can achieve this by implementing a given Compute Optimizer recommendation.

Type: [RDSStorageSavingsOpportunityAfterDiscounts](#) object

Required: No

storageConfiguration

The recommended storage configuration.

Type: [DBStorageConfiguration](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

RDSDBUtilizationMetric

Describes the utilization metric of an Amazon Aurora and RDS database.

To determine the performance difference between your current DB instance and the recommended option, compare the utilization metric data of your service against its projected utilization metric data.

Contents

name

The name of the utilization metric.

Type: String

Valid Values: CPU | Memory | EBSVolumeStorageSpaceUtilization | NetworkReceiveThroughput | NetworkTransmitThroughput | EBSVolumeReadIOPS | EBSVolumeWriteIOPS | EBSVolumeReadThroughput | EBSVolumeWriteThroughput | DatabaseConnections | StorageNetworkReceiveThroughput | StorageNetworkTransmitThroughput | AuroraMemoryHealthState | AuroraMemoryNumDeclinedSql | AuroraMemoryNumKillConnTotal | AuroraMemoryNumKillQueryTotal | ReadIOPSEphemeralStorage | WriteIOPSEphemeralStorage | VolumeReadIOPs | VolumeBytesUsed | VolumeWriteIOPs

Required: No

statistic

The statistic of the utilization metric.

The Compute Optimizer API, AWS Command Line Interface (AWS CLI), and SDKs return utilization metrics using only the Maximum statistic, which is the highest value observed during the specified period.

The Compute Optimizer console displays graphs for some utilization metrics using the Average statistic, which is the value of Sum / SampleCount during the specified period. For more information, see [Viewing resource recommendations](#) in the *AWS Compute Optimizer User Guide*. You can also get averaged utilization metric data for your resources using Amazon CloudWatch. For more information, see the [Amazon CloudWatch User Guide](#).

Type: String

Valid Values: Maximum | Minimum | Average

Required: No

value

The value of the utilization metric.

Type: Double

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

RDSEffectiveRecommendationPreferences

Describes the effective recommendation preferences for Amazon Aurora and RDS databases.

Contents

cpuVendorArchitectures

Describes the CPU vendor and architecture for DB instance recommendations.

Type: Array of strings

Valid Values: AWS_ARM64 | CURRENT

Required: No

enhancedInfrastructureMetrics

Describes the activation status of the enhanced infrastructure metrics preference.

A status of `Active` confirms that the preference is applied in the latest recommendation refresh, and a status of `Inactive` confirms that it's not yet applied to recommendations.

For more information, see [Enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Type: String

Valid Values: `Active` | `Inactive`

Required: No

lookBackPeriod

The number of days the utilization metrics of the DB instance are analyzed.

Type: String

Valid Values: DAYS_14 | DAYS_32 | DAYS_93

Required: No

savingsEstimationMode

Describes the savings estimation mode preference applied for calculating savings opportunity for DB instances.

Type: [RDSavingsEstimationMode](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

RDSInstanceEstimatedMonthlySavings

Describes the estimated monthly savings possible for DB instances by adopting Compute Optimizer recommendations. This is based on DB instance pricing after applying Savings Plans discounts.

Contents

currency

The currency of the estimated monthly savings.

Type: String

Valid Values: USD | CNY

Required: No

value

The value of the estimated monthly savings for DB instances.

Type: Double

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

RDSInstanceSavingsOpportunityAfterDiscounts

Describes the savings opportunity for DB instance recommendations after applying Savings Plans discounts.

Savings opportunity represents the estimated monthly savings after applying Savings Plans discounts. You can achieve this by implementing a given Compute Optimizer recommendation.

Contents

estimatedMonthlySavings

The estimated monthly savings possible by adopting Compute Optimizer's DB instance recommendations. This includes any applicable Savings Plans discounts.

Type: [RDSInstanceEstimatedMonthlySavings](#) object

Required: No

savingsOpportunityPercentage

The estimated monthly savings possible as a percentage of monthly cost by adopting Compute Optimizer's DB instance recommendations. This includes any applicable Savings Plans discounts.

Type: Double

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

RDSSavingsEstimationMode

Describes the savings estimation mode used for calculating savings opportunity for DB instances.

Contents

source

Describes the source for calculating the savings opportunity for DB instances.

Type: String

Valid Values: `PublicPricing` | `CostExplorerRightsizing` | `CostOptimizationHub`

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

RDSStorageEstimatedMonthlySavings

Describes the estimated monthly savings possible for DB instance storage by adopting Compute Optimizer recommendations. This is based on DB instance pricing after applying Savings Plans discounts.

Contents

currency

The currency of the estimated monthly savings.

Type: String

Valid Values: USD | CNY

Required: No

value

The value of the estimated monthly savings for DB instance storage.

Type: Double

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

RDSStorageSavingsOpportunityAfterDiscounts

Describes the savings opportunity for Amazon RDS storage recommendations after applying Savings Plans discounts.

Savings opportunity represents the estimated monthly savings after applying Savings Plans discounts. You can achieve this by implementing a given Compute Optimizer recommendation.

Contents

estimatedMonthlySavings

The estimated monthly savings possible by adopting Compute Optimizer's DB instance storage recommendations. This includes any applicable Savings Plans discounts.

Type: [RDSStorageEstimatedMonthlySavings](#) object

Required: No

savingsOpportunityPercentage

The estimated monthly savings possible as a percentage of monthly cost by adopting Compute Optimizer's DB instance storage recommendations. This includes any applicable Savings Plans discounts.

Type: Double

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

ReasonCodeSummary

A summary of a finding reason code.

Contents

name

The name of the finding reason code.

Type: String

Valid Values: MemoryOverprovisioned | MemoryUnderprovisioned

Required: No

value

The value of the finding reason code summary.

Type: Double

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

RecommendationExportJob

Describes a recommendation export job.

Use the [DescribeRecommendationExportJobs](#) action to view your recommendation export jobs.

Use the [ExportAutoScalingGroupRecommendations](#) or [ExportEC2InstanceRecommendations](#) actions to request an export of your recommendations.

Contents

creationTimestamp

The timestamp of when the export job was created.

Type: Timestamp

Required: No

destination

An object that describes the destination of the export file.

Type: [ExportDestination](#) object

Required: No

failureReason

The reason for an export job failure.

Type: String

Required: No

jobId

The identification number of the export job.

Type: String

Required: No

lastUpdatedTimestamp

The timestamp of when the export job was last updated.

Type: Timestamp

Required: No

resourceType

The resource type of the exported recommendations.

Type: String

Valid Values: Ec2Instance | AutoScalingGroup | EbsVolume | LambdaFunction
| NotApplicable | EcsService | License | RdsDBInstance |
AuroraDBClusterStorage | Idle

Required: No

status

The status of the export job.

Type: String

Valid Values: Queued | InProgress | Complete | Failed

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

RecommendationPreferences

Describes the recommendation preferences to return in the response of a [GetAutoScalingGroupRecommendations](#), [GetEC2InstanceRecommendations](#), [GetEC2RecommendationProjectedMetrics](#), [GetRDSDatabaseRecommendations](#), and [GetRDSDatabaseRecommendationProjectedMetrics](#) request.

Contents

cpuVendorArchitectures

Specifies the CPU vendor and architecture for Amazon EC2 instance and Auto Scaling group recommendations.

For example, when you specify `AWS_ARM64` with:

- A [GetEC2InstanceRecommendations](#) or [GetAutoScalingGroupRecommendations](#) request, Compute Optimizer returns recommendations that consist of Graviton instance types only.
- A [GetEC2RecommendationProjectedMetrics](#) request, Compute Optimizer returns projected utilization metrics for Graviton instance type recommendations only.
- A [ExportEC2InstanceRecommendations](#) or [ExportAutoScalingGroupRecommendations](#) request, Compute Optimizer exports recommendations that consist of Graviton instance types only.

Type: Array of strings

Valid Values: `AWS_ARM64` | `CURRENT`

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

RecommendationPreferencesDetail

Describes a recommendation preference.

Contents

enhancedInfrastructureMetrics

The status of the enhanced infrastructure metrics recommendation preference.

When the recommendations page is refreshed, a status of `Active` confirms that the preference is applied to the recommendations, and a status of `Inactive` confirms that the preference isn't yet applied to recommendations.

For more information, see [Enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Type: String

Valid Values: `Active` | `Inactive`

Required: No

externalMetricsPreference

An object that describes the external metrics recommendation preference.

If the preference is applied in the latest recommendation refresh, an object with a valid `source` value appears in the response. If the preference isn't applied to the recommendations already, then this object doesn't appear in the response.

Type: [ExternalMetricsPreference](#) object

Required: No

inferredWorkloadTypes

The status of the inferred workload types recommendation preference.

When the recommendations page is refreshed, a status of `Active` confirms that the preference is applied to the recommendations, and a status of `Inactive` confirms that the preference isn't yet applied to recommendations.

Type: String

Valid Values: Active | Inactive

Required: No

lookBackPeriod

The preference to control the number of days the utilization metrics of the AWS resource are analyzed. If the preference isn't set, this object is null.

Type: String

Valid Values: DAYS_14 | DAYS_32 | DAYS_93

Required: No

preferredResources

The preference to control which resource type values are considered when generating rightsizing recommendations. This object resolves any wildcard expressions and returns the effective list of candidate resource type values. If the preference isn't set, this object is null.

Type: Array of [EffectivePreferredResource](#) objects

Required: No

resourceType

The target resource type of the recommendation preference to create.

The Ec2Instance option encompasses standalone instances and instances that are part of Auto Scaling groups. The AutoScalingGroup option encompasses only instances that are part of an Auto Scaling group.

Type: String

Valid Values: Ec2Instance | AutoScalingGroup | RdsDBInstance | AuroraDBClusterStorage

Required: No

savingsEstimationMode

Describes the savings estimation mode used for calculating savings opportunity.

Only the account manager or delegated administrator of your organization can activate this preference.

Type: String

Valid Values: `AfterDiscounts` | `BeforeDiscounts`

Required: No

scope

An object that describes the scope of the recommendation preference.

Recommendation preferences can be created at the organization level (for management accounts of an organization only), account level, and resource level. For more information, see [Activating enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Type: [Scope](#) object

Required: No

utilizationPreferences

The preference to control the resource's CPU utilization threshold, CPU utilization headroom, and memory utilization headroom. If the preference isn't set, this object is null.

Note

This preference is only available for the Amazon EC2 instance resource type.

Type: Array of [UtilizationPreference](#) 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 Java V2](#)

- [AWS SDK for Ruby V3](#)

RecommendationSource

Describes the source of a recommendation, such as an Amazon EC2 instance or Auto Scaling group.

Contents

recommendationSourceArn

The Amazon Resource Name (ARN) of the recommendation source.

Type: String

Required: No

recommendationSourceType

The resource type of the recommendation source.

Type: String

Valid Values: Ec2Instance | AutoScalingGroup | EbsVolume | LambdaFunction
| EcsService | License | RdsDBInstance | RdsDBInstanceStorage |
AuroraDBClusterStorage

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

RecommendationSummary

A summary of a recommendation.

Contents

accountId

The AWS account ID of the recommendation summary.

Type: String

Required: No

aggregatedSavingsOpportunity

Describes the savings opportunity for recommendations of a given resource type or for the recommendation option of an individual resource.

Savings opportunity represents the estimated monthly savings you can achieve by implementing a given Compute Optimizer recommendation.

Important

Savings opportunity data requires that you opt in to Cost Explorer, as well as activate **Receive Amazon EC2 resource recommendations** in the Cost Explorer preferences page. That creates a connection between Cost Explorer and Compute Optimizer. With this connection, Cost Explorer generates savings estimates considering the price of existing resources, the price of recommended resources, and historical usage data. Estimated monthly savings reflects the projected dollar savings associated with each of the recommendations generated. For more information, see [Enabling Cost Explorer](#) and [Optimizing your cost with Rightsizing Recommendations](#) in the *Cost Management User Guide*.

Type: [SavingsOpportunity](#) object

Required: No

currentPerformanceRiskRatings

An object that describes the performance risk ratings for a given resource type.

Type: [CurrentPerformanceRiskRatings](#) object

Required: No

idleSavingsOpportunity

Describes the savings opportunity for recommendations of a given resource type or for the recommendation option of an individual resource.

Savings opportunity represents the estimated monthly savings you can achieve by implementing a given Compute Optimizer recommendation.

Important

Savings opportunity data requires that you opt in to Cost Explorer, as well as activate **Receive Amazon EC2 resource recommendations** in the Cost Explorer preferences page. That creates a connection between Cost Explorer and Compute Optimizer. With this connection, Cost Explorer generates savings estimates considering the price of existing resources, the price of recommended resources, and historical usage data. Estimated monthly savings reflects the projected dollar savings associated with each of the recommendations generated. For more information, see [Enabling Cost Explorer](#) and [Optimizing your cost with Rightsizing Recommendations](#) in the *Cost Management User Guide*.

Type: [SavingsOpportunity](#) object

Required: No

idleSummaries

Describes the findings summary of the idle resources.

Type: Array of [IdleSummary](#) objects

Required: No

inferredWorkloadSavings

An array of objects that describes the estimated monthly saving amounts for the instances running on the specified `inferredWorkloadTypes`. The array contains the top five savings opportunities for the instances that run inferred workload types.

Type: Array of [InferredWorkloadSaving](#) objects

Required: No

recommendationResourceType

The resource type that the recommendation summary applies to.

Type: String

Valid Values: Ec2Instance | AutoScalingGroup | EbsVolume | LambdaFunction
| EcsService | License | RdsDBInstance | RdsDBInstanceStorage |
AuroraDBClusterStorage

Required: No

savingsOpportunity

An object that describes the savings opportunity for a given resource type. Savings opportunity includes the estimated monthly savings amount and percentage.

Type: [SavingsOpportunity](#) object

Required: No

summaries

An array of objects that describe a recommendation summary.

Type: Array of [Summary](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

RecommendedOptionProjectedMetric

Describes a projected utilization metric of a recommendation option.

Note

The Cpu and Memory metrics are the only projected utilization metrics returned when you run the [GetEC2RecommendationProjectedMetrics](#) action. Additionally, the Memory metric is returned only for resources that have the unified CloudWatch agent installed on them. For more information, see [Enabling Memory Utilization with the CloudWatch Agent](#).

Contents

projectedMetrics

An array of objects that describe a projected utilization metric.

Type: Array of [ProjectedMetric](#) objects

Required: No

rank

The rank of the recommendation option projected metric.

The top recommendation option is ranked as 1.

The projected metric rank correlates to the recommendation option rank. For example, the projected metric ranked as 1 is related to the recommendation option that is also ranked as 1 in the same response.

Type: Integer

Required: No

recommendedInstanceType

The recommended instance type.

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

S3Destination

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and object keys of a recommendations export file, and its associated metadata file.

Contents

bucket

The name of the Amazon S3 bucket used as the destination of an export file.

Type: String

Required: No

key

The Amazon S3 bucket key of an export file.

The key uniquely identifies the object, or export file, in the S3 bucket.

Type: String

Required: No

metadataKey

The Amazon S3 bucket key of a metadata file.

The key uniquely identifies the object, or metadata file, in the S3 bucket.

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 Java V2](#)

- [AWS SDK for Ruby V3](#)

S3DestinationConfig

Describes the destination Amazon Simple Storage Service (Amazon S3) bucket name and key prefix for a recommendations export job.

You must create the destination Amazon S3 bucket for your recommendations export before you create the export job. Compute Optimizer does not create the S3 bucket for you. After you create the S3 bucket, ensure that it has the required permission policy to allow Compute Optimizer to write the export file to it. If you plan to specify an object prefix when you create the export job, you must include the object prefix in the policy that you add to the S3 bucket. For more information, see [Amazon S3 Bucket Policy for Compute Optimizer](#) in the *Compute Optimizer User Guide*.

Contents

bucket

The name of the Amazon S3 bucket to use as the destination for an export job.

Type: String

Required: No

keyPrefix

The Amazon S3 bucket prefix for an export job.

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

SavingsOpportunity

Describes the savings opportunity for recommendations of a given resource type or for the recommendation option of an individual resource.

Savings opportunity represents the estimated monthly savings you can achieve by implementing a given Compute Optimizer recommendation.

Important

Savings opportunity data requires that you opt in to Cost Explorer, as well as activate **Receive Amazon EC2 resource recommendations** in the Cost Explorer preferences page. That creates a connection between Cost Explorer and Compute Optimizer. With this connection, Cost Explorer generates savings estimates considering the price of existing resources, the price of recommended resources, and historical usage data. Estimated monthly savings reflects the projected dollar savings associated with each of the recommendations generated. For more information, see [Enabling Cost Explorer](#) and [Optimizing your cost with Rightsizing Recommendations](#) in the *Cost Management User Guide*.

Contents

estimatedMonthlySavings

An object that describes the estimated monthly savings amount possible by adopting Compute Optimizer recommendations for a given resource. This is based on the On-Demand instance pricing..

Type: [EstimatedMonthlySavings](#) object

Required: No

savingsOpportunityPercentage

The estimated monthly savings possible as a percentage of monthly cost by adopting Compute Optimizer recommendations for a given resource.

Type: Double

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

Scope

Describes the scope of a recommendation preference.

Recommendation preferences can be created at the organization level (for management accounts of an organization only), account level, and resource level. For more information, see [Activating enhanced infrastructure metrics](#) in the *AWS Compute Optimizer User Guide*.

Note

You cannot create recommendation preferences for Auto Scaling groups at the organization and account levels. You can create recommendation preferences for Auto Scaling groups only at the resource level by specifying a scope name of `ResourceArn` and a scope value of the Auto Scaling group Amazon Resource Name (ARN). This will configure the preference for all instances that are part of the specified Auto Scaling group. You also cannot create recommendation preferences at the resource level for instances that are part of an Auto Scaling group. You can create recommendation preferences at the resource level only for standalone instances.

Contents

name

The name of the scope.

The following scopes are possible:

- `Organization` - Specifies that the recommendation preference applies at the organization level, for all member accounts of an organization.
- `AccountId` - Specifies that the recommendation preference applies at the account level, for all resources of a given resource type in an account.
- `ResourceArn` - Specifies that the recommendation preference applies at the individual resource level.

Type: String

Valid Values: `Organization` | `AccountId` | `ResourceArn`

Required: No

value

The value of the scope.

If you specified the name of the scope as:

- `Organization` - The value must be `ALL_ACCOUNTS`.
- `AccountId` - The value must be a 12-digit AWS account ID.
- `ResourceArn` - The value must be the Amazon Resource Name (ARN) of an EC2 instance or an Auto Scaling group.

Only EC2 instance and Auto Scaling group ARNs are currently supported.

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

ServiceConfiguration

The Amazon ECS service configurations used for recommendations.

Contents

autoScalingConfiguration

Describes the Auto Scaling configuration methods for an Amazon ECS service. This affects the generated recommendations. For example, if Auto Scaling is configured on a service's CPU, then Compute Optimizer doesn't generate CPU size recommendations.

The Auto Scaling configuration methods include:

- `TARGET_TRACKING_SCALING_CPU` — If the Amazon ECS service is configured to use target scaling on CPU, Compute Optimizer doesn't generate CPU recommendations.
- `TARGET_TRACKING_SCALING_MEMORY` — If the Amazon ECS service is configured to use target scaling on memory, Compute Optimizer doesn't generate memory recommendations.

For more information about step scaling and target scaling, see [Step scaling policies for Application Auto Scaling](#) and [Target tracking scaling policies for Application Auto Scaling](#) in the *Application Auto Scaling User Guide*.

Type: String

Valid Values: `TargetTrackingScalingCpu` | `TargetTrackingScalingMemory`

Required: No

containerConfigurations

The container configurations within a task of an Amazon ECS service.

Type: Array of [ContainerConfiguration](#) objects

Required: No

cpu

The number of CPU units used by the tasks in the Amazon ECS service.

Type: Integer

Required: No

memory

The amount of memory used by the tasks in the Amazon ECS service.

Type: Integer

Required: No

taskDefinitionArn

The task definition ARN used by the tasks in the Amazon ECS service.

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

Summary

The summary of a recommendation.

Contents

name

The finding classification of the recommendation.

Type: String

Valid Values: Underprovisioned | Overprovisioned | Optimized | NotOptimized

Required: No

reasonCodeSummaries

An array of objects that summarize a finding reason code.

Type: Array of [ReasonCodeSummary](#) objects

Required: No

value

The value of the recommendation summary.

Type: Double

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

Tag

A list of tag key and value pairs that you define.

Contents

key

One part of a key-value pair that makes up a tag. A key is a general label that acts like a category for more specific tag values.

Type: String

Required: No

value

One part of a key-value pair that make up a tag. A value acts as a descriptor within a tag category (key). The value can be empty or null.

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

UtilizationMetric

Describes a utilization metric of a resource, such as an Amazon EC2 instance.

Compare the utilization metric data of your resource against its projected utilization metric data to determine the performance difference between your current resource and the recommended option.

Contents

name

The name of the utilization metric.

The following utilization metrics are available:

- **Cpu** - The percentage of allocated EC2 compute units that are currently in use on the instance. This metric identifies the processing power required to run an application on the instance.

Depending on the instance type, tools in your operating system can show a lower percentage than CloudWatch when the instance is not allocated a full processor core.

Units: Percent

- **Memory** - The percentage of memory that is currently in use on the instance. This metric identifies the amount of memory required to run an application on the instance.

Units: Percent

Note

The Memory metric is returned only for resources that have the unified CloudWatch agent installed on them. For more information, see [Enabling Memory Utilization with the CloudWatch Agent](#).

- **GPU** - The percentage of allocated GPUs that currently run on the instance.
- **GPU_MEMORY** - The percentage of total GPU memory that currently runs on the instance.

Note

The GPU and GPU_MEMORY metrics are only returned for resources with the unified CloudWatch Agent installed on them. For more information, see [Enabling NVIDIA GPU utilization with the CloudWatch Agent](#).

- **EBS_READ_OPS_PER_SECOND** - The completed read operations from all EBS volumes attached to the instance in a specified period of time.

Unit: Count

- **EBS_WRITE_OPS_PER_SECOND** - The completed write operations to all EBS volumes attached to the instance in a specified period of time.

Unit: Count

- **EBS_READ_BYTES_PER_SECOND** - The bytes read from all EBS volumes attached to the instance in a specified period of time.

Unit: Bytes

- **EBS_WRITE_BYTES_PER_SECOND** - The bytes written to all EBS volumes attached to the instance in a specified period of time.

Unit: Bytes

- **DISK_READ_OPS_PER_SECOND** - The completed read operations from all instance store volumes available to the instance in a specified period of time.

If there are no instance store volumes, either the value is 0 or the metric is not reported.

- **DISK_WRITE_OPS_PER_SECOND** - The completed write operations from all instance store volumes available to the instance in a specified period of time.

If there are no instance store volumes, either the value is 0 or the metric is not reported.

- **DISK_READ_BYTES_PER_SECOND** - The bytes read from all instance store volumes available to the instance. This metric is used to determine the volume of the data the application reads from the disk of the instance. This can be used to determine the speed of the application.

If there are no instance store volumes, either the value is 0 or the metric is not reported.

- **DISK_WRITE_BYTES_PER_SECOND** - The bytes written to all instance store volumes available to the instance. This metric is used to determine the volume of the data the

application writes onto the disk of the instance. This can be used to determine the speed of the application.

If there are no instance store volumes, either the value is 0 or the metric is not reported.

- **NETWORK_IN_BYTES_PER_SECOND** - The number of bytes received by the instance on all network interfaces. This metric identifies the volume of incoming network traffic to a single instance.
- **NETWORK_OUT_BYTES_PER_SECOND** - The number of bytes sent out by the instance on all network interfaces. This metric identifies the volume of outgoing network traffic from a single instance.
- **NETWORK_PACKETS_IN_PER_SECOND** - The number of packets received by the instance on all network interfaces. This metric identifies the volume of incoming traffic in terms of the number of packets on a single instance.
- **NETWORK_PACKETS_OUT_PER_SECOND** - The number of packets sent out by the instance on all network interfaces. This metric identifies the volume of outgoing traffic in terms of the number of packets on a single instance.

Type: String

Valid Values: Cpu | Memory | EBS_READ_OPS_PER_SECOND | EBS_WRITE_OPS_PER_SECOND | EBS_READ_BYTES_PER_SECOND | EBS_WRITE_BYTES_PER_SECOND | DISK_READ_OPS_PER_SECOND | DISK_WRITE_OPS_PER_SECOND | DISK_READ_BYTES_PER_SECOND | DISK_WRITE_BYTES_PER_SECOND | NETWORK_IN_BYTES_PER_SECOND | NETWORK_OUT_BYTES_PER_SECOND | NETWORK_PACKETS_IN_PER_SECOND | NETWORK_PACKETS_OUT_PER_SECOND | GPU_PERCENTAGE | GPU_MEMORY_PERCENTAGE

Required: No

statistic

The statistic of the utilization metric.

The Compute Optimizer API, AWS Command Line Interface (AWS CLI), and SDKs return utilization metrics using only the Maximum statistic, which is the highest value observed during the specified period.

The Compute Optimizer console displays graphs for some utilization metrics using the Average statistic, which is the value of Sum / SampleCount during the specified period. For more

information, see [Viewing resource recommendations](#) in the *AWS Compute Optimizer User Guide*. You can also get averaged utilization metric data for your resources using Amazon CloudWatch. For more information, see the [Amazon CloudWatch User Guide](#).

Type: String

Valid Values: Maximum | Average

Required: No

value

The value of the utilization metric.

Type: Double

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

UtilizationPreference

The preference to control the resource's CPU utilization threshold, CPU utilization headroom, and memory utilization headroom.

Note

This preference is only available for the Amazon EC2 instance resource type.

Contents

metricName

The name of the resource utilization metric name to customize.

Type: String

Valid Values: `CpuUtilization` | `MemoryUtilization`

Required: No

metricParameters

The parameters to set when customizing the resource utilization thresholds.

Type: [CustomizableMetricParameters](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

VolumeConfiguration

Describes the configuration of an Amazon Elastic Block Store (Amazon EBS) volume.

Contents

rootVolume

Contains the image used to boot the instance during launch.

Type: Boolean

Required: No

volumeBaselineIOPS

The baseline IOPS of the volume.

Type: Integer

Required: No

volumeBaselineThroughput

The baseline throughput of the volume.

Type: Integer

Required: No

volumeBurstIOPS

The burst IOPS of the volume.

Type: Integer

Required: No

volumeBurstThroughput

The burst throughput of the volume.

Type: Integer

Required: No

volumeSize

The size of the volume, in GiB.

Type: Integer

Required: No

volumeType

The volume type.

The volume types can be the following:

- General Purpose SSD gp2 and gp3
- Provisioned IOPS SSD io1, io2, and io2 Block Express
- Throughput Optimized HDD st1
- Cold HDD sc1
- Magnetic volumes standard

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

VolumeRecommendation

Describes an Amazon Elastic Block Store (Amazon EBS) volume recommendation.

Contents

accountId

The AWS account ID of the volume.

Type: String

Required: No

currentConfiguration

An array of objects that describe the current configuration of the volume.

Type: [VolumeConfiguration](#) object

Required: No

currentPerformanceRisk

The risk of the current EBS volume not meeting the performance needs of its workloads. The higher the risk, the more likely the current EBS volume doesn't have sufficient capacity.

Type: String

Valid Values: VeryLow | Low | Medium | High

Required: No

effectiveRecommendationPreferences

Describes the effective recommendation preferences for Amazon EBS volume.

Type: [EBSEffectiveRecommendationPreferences](#) object

Required: No

finding

The finding classification of the volume.

Findings for volumes include:

- **NotOptimized** —A volume is considered not optimized when AWS Compute Optimizer identifies a recommendation that can provide better performance for your workload.
- **Optimized** —An volume is considered optimized when Compute Optimizer determines that the volume is correctly provisioned to run your workload based on the chosen volume type. For optimized resources, Compute Optimizer might recommend a new generation volume type.

Type: String

Valid Values: Optimized | NotOptimized

Required: No

lastRefreshTimestamp

The timestamp of when the volume recommendation was last generated.

Type: Timestamp

Required: No

lookBackPeriodInDays

The number of days for which utilization metrics were analyzed for the volume.

Type: Double

Required: No

tags

A list of tags assigned to your Amazon EBS volume recommendations.

Type: Array of [Tag](#) objects

Required: No

utilizationMetrics

An array of objects that describe the utilization metrics of the volume.

Type: Array of [EBSUtilizationMetric](#) objects

Required: No

volumeArn

The Amazon Resource Name (ARN) of the current volume.

Type: String

Required: No

volumeRecommendationOptions

An array of objects that describe the recommendation options for the volume.

Type: Array of [VolumeRecommendationOption](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

VolumeRecommendationOption

Describes a recommendation option for an Amazon Elastic Block Store (Amazon EBS) instance.

Contents

configuration

An array of objects that describe a volume configuration.

Type: [VolumeConfiguration](#) object

Required: No

performanceRisk

The performance risk of the volume recommendation option.

Performance risk is the likelihood of the recommended volume type meeting the performance requirement of your workload.

The value ranges from 0 - 4, with 0 meaning that the recommended resource is predicted to always provide enough hardware capability. The higher the performance risk is, the more likely you should validate whether the recommendation will meet the performance requirements of your workload before migrating your resource.

Type: Double

Valid Range: Minimum value of 0. Maximum value of 4.

Required: No

rank

The rank of the volume recommendation option.

The top recommendation option is ranked as 1.

Type: Integer

Required: No

savingsOpportunity

An object that describes the savings opportunity for the EBS volume recommendation option. Savings opportunity includes the estimated monthly savings amount and percentage.

Type: [SavingsOpportunity](#) object

Required: No

savingsOpportunityAfterDiscounts

An object that describes the savings opportunity for the Amazon EBS volume recommendation option with specific discounts. Savings opportunity includes the estimated monthly savings and percentage.

Type: [EBSSavingsOpportunityAfterDiscounts](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

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 [Signing AWS API requests](#) in the *IAM User Guide*.

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 [Create a signed AWS API request](#) in the *IAM User Guide*.

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'T'HHMMSS'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 [Elements of an AWS API request signature](#) in the *IAM User Guide*.

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 STS, see [AWS services that work with IAM](#) in the *IAM User Guide*.

Condition: If you're using temporary security credentials from AWS STS, 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 [Create a signed AWS API request](#) in the *IAM User Guide*.

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

Type: string

Required: Conditional

Common Errors

This section lists the errors common to the API actions of all AWS services. For errors specific to an API action for this service, see the topic for that API action.

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

IncompleteSignature

The request signature does not conform to AWS standards.

HTTP Status Code: 400

InternalFailure

The request processing has failed because of an unknown error, exception or failure.

HTTP Status Code: 500

InvalidAction

The action or operation requested is invalid. Verify that the action is typed correctly.

HTTP Status Code: 400

InvalidClientTokenId

The X.509 certificate or AWS access key ID provided does not exist in our records.

HTTP Status Code: 403

NotAuthorized

You do not have permission to perform this action.

HTTP Status Code: 400

OptInRequired

The AWS access key ID needs a subscription for the service.

HTTP Status Code: 403

RequestExpired

The request reached the service more than 15 minutes after the date stamp on the request or more than 15 minutes after the request expiration date (such as for pre-signed URLs), or the date stamp on the request is more than 15 minutes in the future.

HTTP Status Code: 400

ServiceUnavailable

The request has failed due to a temporary failure of the server.

HTTP Status Code: 503

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

ValidationError

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400