

---

# AWS Compute Optimizer

## User Guide



## **AWS Compute Optimizer: User Guide**

Copyright © 2020 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

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

## Table of Contents

What is Compute Optimizer? .....	1
Supported resources and requirements .....	1
Opting in .....	1
Metrics analyzed .....	1
Findings and recommendations .....	1
Availability .....	2
Requirements .....	3
Metric requirements .....	3
Amazon EC2 instance requirements .....	3
Auto Scaling group requirements .....	3
Getting started .....	4
Accounts supported by Compute Optimizer .....	4
Opting in your account .....	4
Opting out your account .....	5
Controlling access .....	5
Granting access for standalone AWS accounts .....	6
Granting access for a management account of an organization .....	7
Denying access .....	8
Using Service-Linked Roles .....	8
Service-Linked Role permissions for Compute Optimizer .....	8
Service-Linked Role permissions .....	9
Creating a Service-Linked Role for Compute Optimizer .....	9
Editing a Service-Linked Role for Compute Optimizer .....	10
Deleting a Service-Linked Role for Compute Optimizer .....	10
Supported Regions for Compute Optimizer Service-Linked Roles .....	10
S3 bucket policy for Compute Optimizer .....	10
Using encrypted S3 buckets for your recommendations export .....	11
Specifying an existing bucket for your recommendations export .....	11
Additional resources .....	13
Metrics analyzed .....	14
Enabling memory utilization with the CloudWatch Agent .....	14
Enabling detailed monitoring in CloudWatch .....	15
Viewing the dashboard .....	16
Findings classifications .....	16
EC2 instances .....	16
Auto Scaling groups .....	16
Viewing the dashboard .....	17
Viewing resource recommendations .....	18
Supported resources and requirements .....	18
Findings classifications .....	18
Prices and purchasing options .....	18
Estimated monthly savings .....	19
Viewing recommendations .....	19
Viewing resource details .....	21
Performance risk .....	21
Prices and purchasing options on the resource details page .....	21
Utilization graphs .....	21
Viewing resource details .....	22
Exporting recommendations .....	24
Restrictions .....	24
Amazon S3 bucket permission requirements .....	24
Exporting your recommendations .....	24
Viewing your export jobs .....	26
Exported files .....	26

Recommendations file .....	26
Metadata file .....	36
Document history .....	39

# What is AWS Compute Optimizer?

AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of your AWS resources, such as Amazon Elastic Compute Cloud (Amazon EC2) instances and Auto Scaling groups. 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 graphs showing recent utilization metric history data, as well as projected utilization for recommendations, which you can use to evaluate which recommendation provides the best price-performance trade-off. The analysis and visualization of your usage patterns can help you decide when to move or resize your running resources, and still meet your performance and capacity requirements.

Compute Optimizer provides a [console experience](#), and a [set of APIs](#) that allows you to view the findings and recommendations for your resources across multiple AWS Regions. You can also view findings and recommendations across multiple accounts, if you opt in the management account of an organization. The findings from the service are also reported in the consoles of the supported services, such as the Amazon EC2 console.

## Supported resources and requirements

Compute Optimizer generates recommendations for Amazon Elastic Compute Cloud (Amazon EC2) instances and Amazon EC2 Auto Scaling groups that meet a specific set of requirements, and that have accumulated sufficient metric data. For more information, see [Supported resources and requirements](#) (p. 3).

## Opting in

You must opt in to have Compute Optimizer analyze your AWS resources. The service supports standalone AWS accounts, member accounts of an organization, and the management account of an organization. For more information, see [Getting started with AWS Compute Optimizer](#) (p. 4).

## Metrics analyzed

After you opt in, Compute Optimizer begins analyzing the specifications (vCPUs, memory, storage, and so on) and the utilization metrics of your resources from Amazon CloudWatch. For more information, see [Metrics analyzed by AWS Compute Optimizer](#) (p. 14).

## Findings and recommendations

Optimization findings for your resources are displayed on the Compute Optimizer dashboard. For more information, see [Viewing the AWS Compute Optimizer dashboard](#) (p. 16).

The top optimization recommendations for each of your resources are listed on the recommendations page. For more information, see [Viewing resource recommendations](#) (p. 18).

Additional optimization recommendations and utilization graphs for a specific resource are listed on the resource details page. For more information, see [Viewing resource details](#) (p. 21).

Export your optimization recommendations to record them over time, and share the data with others. For more information, see [Exporting recommendations \(p. 24\)](#).

## Availability

To view the currently supported AWS Regions and endpoints for Compute Optimizer, see [Compute Optimizer Endpoints and Quotas](#) in the *AWS General Reference*.

# Supported resources and requirements

AWS Compute Optimizer generates recommendations for Amazon Elastic Compute Cloud (Amazon EC2) instances, and Amazon EC2 Auto Scaling groups that meet the following configuration and metric requirements.

## Metric requirements

Compute Optimizer requires at least 30 *consecutive* hours of metric data from your resource to generate recommendations. Wait and allow more time for metric data to accumulate if your resources meet the above requirements but they are not listed in the Compute Optimizer console. For more information about the metrics that are analyzed, see [Metrics analyzed by AWS Compute Optimizer \(p. 14\)](#).

After the analysis is completed, which could take up to 12 hours, Compute Optimizer presents its findings on the dashboard page. For more information, see [Viewing the AWS Compute Optimizer dashboard \(p. 16\)](#).

## Amazon EC2 instance requirements

Compute Optimizer generates recommendations for M, C, R, T, and X instance types. However, it does not generate recommendations for -g instance types powered by AWS Graviton2 processors (e.g., C6g), and for -n instance types that have higher network bandwidth performance (e.g., M5n).

If you're running instances that use unsupported instance types, in addition to those supported by Compute Optimizer, the service generates recommendations only for the supported instances. To determine the supported Regions for each of these instance types, see [Finding an Amazon EC2 instance type](#) in the *Amazon EC2 User Guide for Linux Instances*.

## Auto Scaling group requirements

Compute Optimizer generates recommendations for instances in Auto Scaling groups that are configured to launch and run M, C, R, T, and X instance types. However, it does not generate recommendations for -g instance types powered by AWS Graviton2 processors (e.g., C6g), and for -n instance types that have higher network bandwidth performance (e.g., M5n).

The Auto Scaling groups must also be configured to run a single instance type (i.e., no mixed instance types), must not have a scaling policy attached to them, and have the same values for desired, minimum, and maximum capacity (i.e., an Auto Scaling group with a fixed number of instances). Compute Optimizer generates recommendations for instances in Auto Scaling groups that meet *all* of these configuration requirements.

# Getting started with AWS Compute Optimizer

When you access the AWS Compute Optimizer console for the first time, you are asked to opt in, using the account that you're signed in with, before you can use the service. You can also opt in, and opt out using the Compute Optimizer API, AWS Command Line Interface (AWS CLI), or SDKs. For more information, see [UpdateEnrollmentStatus](#) in the *AWS Compute Optimizer API Reference*.

By opting in, you are authorizing Compute Optimizer to analyze the specifications and utilization metrics of your AWS resources, such as EC2 instances and Auto Scaling groups. For more information, see [Metrics analyzed by AWS Compute Optimizer \(p. 14\)](#). To access the data, Compute Optimizer automatically creates a Service-Linked Role in your account. For more information, see [Using Service-Linked Roles for AWS Compute Optimizer \(p. 8\)](#).

## Accounts supported by Compute Optimizer

The following AWS account types can opt in to Compute Optimizer:

- **Standalone AWS account** - A standalone AWS account that does not have AWS Organizations enabled. If you opt in to Compute Optimizer while signed in to a standalone account, the service analyzes resources that are in the account, and generates optimization recommendations for those resources.
- **Member account of an organization** - An AWS account that is a member of an organization. If you opt in to Compute Optimizer while signed in to a member account of an organization, the service analyzes resources that are in this member account, and generates optimization recommendations for those resources.
- **Management account of an organization** - An AWS account that administers an organization. If you opt in to Compute Optimizer while signed in to a management account of an organization, the service gives you the option to opt in only the management account, or all accounts in the organization.

When you opt in all accounts in the organization, Compute Optimizer analyzes resources that are in all member accounts of the organization. You can then view recommendations for resources in each member account.

### Important

To successfully opt in all accounts in an organization, the organization must have all features enabled. For more information, see [Enabling All Features in Your Organization](#) in the *AWS Organizations User Guide*.

## Opting in your account

Use the following procedure to opt in your account.

### To opt in your account

1. Open the Compute Optimizer console at <https://console.aws.amazon.com/compute-optimizer/>.  
If this is your first time using the Compute Optimizer console, the **Compute Optimizer landing page** is displayed.
2. Choose **Get started**.
3. On the **Account setup** page, review the **Getting started** and **Setting up your account** sections.



- The following options are displayed if the account that you're signed in to is a management account of an organization. Choose one before continuing to the next step.
  - Only this account** - Choose this option to opt in only the account that you're currently signed in to. If you choose this option, Compute Optimizer analyzes resources that are in the individual account, and generates optimization recommendations for those resources.
  - All accounts within this organization** - Choose this option to opt in the account you're currently signed in to, and all of its member accounts. If you choose this option, Compute Optimizer analyzes resources that are in all accounts in the organization, and generates optimization recommendations for those resources.
- Choose **Opt in**. By opting in, you indicate that you agree to and understand the requirements to opt in to Compute Optimizer.

After you opt in, you are redirected to the dashboard in the Compute Optimizer console. The service begins analyzing the configuration and utilization metrics of your AWS resources. However, the findings and optimization recommendations might not be generated right away. For more information, see [Viewing the AWS Compute Optimizer dashboard \(p. 16\)](#).

**Note**

To improve the recommendation quality of Compute Optimizer, AWS may use your CloudWatch metrics and configuration data, such as memory utilization, to improve our recommendation models and algorithms. If you want to opt out of this experience and request that AWS does not use your CloudWatch metrics or configuration data to improve recommendation quality of Compute Optimizer, contact AWS Support.

## Opting out your account

Use the following procedure to opt out your account from Compute Optimizer using the command line.

**To opt out an account**

- Open a Terminal or Command Prompt window.

If you haven't already, install the AWS CLI and configure it to work with Compute Optimizer. For more information, see [Installing the AWS CLI](#) and [Quickly Configuring the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

- Enter the following command.

```
aws compute-optimizer update-enrollment-status --status Inactive
```

**Note**

You cannot specify the `--include-member-accounts` parameter when opting out with the above command; you will receive an error message.

Your account is opted out of Compute Optimizer after running the previous command. If you access the Compute Optimizer console, you should see the option to opt in again.

## Controlling access with AWS Identity and Access Management

You can use AWS Identity and Access Management (IAM) to create identities (users, groups, or roles), and then give those identities permissions to access the AWS Compute Optimizer console and APIs.

By default, IAM users do not have access to the Compute Optimizer console and APIs. You give users access by attaching IAM policies to a single user, a group of users, or a role. For more information, see [Identities \(Users, Groups, and Roles\)](#) and [Overview of IAM Policies in the IAM User Guide](#).

After you create IAM users, you can give those users individual passwords. Then, they can sign in to your account and view Compute Optimizer information by using an account-specific sign-in page. For more information, see [How Users Sign In to Your Account](#).

**Important**

In order to view recommendations for EC2 instances, an IAM user must have `ec2:DescribeInstances` permission. In order to view recommendations for Auto Scaling groups, an IAM user must have `autoscaling:DescribeAutoScalingGroups` permission. In order to view current CloudWatch metrics data in the Compute Optimizer console, an IAM user must have `cloudwatch:GetMetricData` permissions.

If the user or group that you want to give permissions to already has a policy, you can add one of the Compute Optimizer-specific policy statements illustrated here to that policy.

## Granting access for standalone AWS accounts

The following policy statement grants full access to Compute Optimizer for standalone AWS accounts.

**Note**

To get access to Compute Optimizer you must also add a policy that grants the ability to create a service-linked role. For more information, see [Using Service-Linked Roles for AWS Compute Optimizer \(p. 8\)](#).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "compute-optimizer:*",
        "ec2:DescribeInstances",
        "autoscaling:DescribeAutoScalingGroups",
        "cloudwatch:GetMetricData"
      ],
      "Resource": "*"
    }
  ]
}
```

The following policy statement grants read-only access to Compute Optimizer for standalone AWS accounts.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "compute-optimizer:GetEnrollmentStatus",
        "compute-optimizer:GetRecommendationSummaries",
        "compute-optimizer:GetEC2InstanceRecommendations",
        "compute-optimizer:GetEC2RecommendationProjectedMetrics",
        "compute-optimizer:GetAutoScalingGroupRecommendations",
        "compute-optimizer:DescribeRecommendationExportJobs",
        "ec2:DescribeInstances",
        "autoscaling:DescribeAutoScalingGroups",
      ]
    }
  ]
}
```

```
        "cloudwatch:GetMetricData"  
      ],  
      "Resource": "*"   
    }  
  ]  
}
```

## Granting access for a management account of an organization

The following policy statement grants full access to Compute Optimizer for a management account of an organization.

### Note

To get access to Compute Optimizer you must also add a policy that grants the ability to create a service-linked role. For more information, see [Using Service-Linked Roles for AWS Compute Optimizer \(p. 8\)](#).

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "compute-optimizer:*",  
        "ec2:DescribeInstances",  
        "autoscaling:DescribeAutoScalingGroups",  
        "cloudwatch:GetMetricData",  
        "organizations:ListAccounts",  
        "organizations:DescribeOrganization",  
        "organizations:DescribeAccount",  
        "organizations:EnableAWSServiceAccess",  
      ],  
      "Resource": "*"   
    }  
  ]  
}
```

The following policy statement grants read-only access to Compute Optimizer for a management account of an organization.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "compute-optimizer:GetEnrollmentStatus",  
        "compute-optimizer:GetRecommendationSummaries",  
        "compute-optimizer:GetEC2InstanceRecommendations",  
        "compute-optimizer:GetEC2RecommendationProjectedMetrics",  
        "compute-optimizer:GetAutoScalingGroupRecommendations",  
        "ec2:DescribeInstances",  
        "autoscaling:DescribeAutoScalingGroups",  
        "cloudwatch:GetMetricData",  
        "organizations:ListAccounts",  
        "organizations:DescribeOrganization",  
        "organizations:DescribeAccount"  
      ],  
      "Resource": "*"   
    }  
  ]  
}
```

```
}  
]  
}
```

## Denying access

The following policy statement denies access to Compute Optimizer.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Deny",  
      "Action": "compute-optimizer:*",  
      "Resource": "*"  
    }  
  ]  
}
```

# Using Service-Linked Roles for AWS Compute Optimizer

AWS Compute Optimizer uses AWS Identity and Access Management (IAM) [service-linked roles](#). A service-linked role is a unique type of IAM role that is linked directly to Compute Optimizer. Service-linked roles are predefined by Compute Optimizer and include all of the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes setting up Compute Optimizer easier because you don't have to manually add the necessary permissions. Compute Optimizer defines the permissions of its service-linked roles, and unless defined otherwise, only Compute Optimizer can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

For information about other services that support service-linked roles, see [AWS Services That Work with IAM](#) and look for the services that have **Yes** in the **Service-Linked Role** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

## Service-Linked Role permissions for Compute Optimizer

Compute Optimizer uses the service-linked role named **AWSServiceRoleForComputeOptimizer** – Role to access Amazon CloudWatch metrics for AWS resources in the account.

The **AWSServiceRoleForComputeOptimizer** service-linked role trusts the following services to assume the role:

- `compute-optimizer.amazonaws.com`

The role permissions policy allows Compute Optimizer to complete the following actions on the specified resources:

- Action: `cloudwatch:GetMetricData` on all AWS resources.

- Action: `organizations:DescribeOrganization` on all AWS resources.
- Action: `organizations:ListAccounts` on all AWS resources.
- Action: `organizations:ListAWSServiceAccessForOrganization` on all AWS resources.

## Service-Linked Role permissions

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create a service-linked role for Compute Optimizer. For more information, see [Service-Linked Role Permissions](#) in the *IAM User Guide*.

### To allow an IAM entity to create a specific service-linked role for Compute Optimizer

Add the following policy to the IAM entity that needs to create the service-linked role.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iam:CreateServiceLinkedRole",
      "Resource": "arn:aws:iam::*:role/aws-service-role/compute-optimizer.amazonaws.com/AWSServiceRoleForComputeOptimizer*",
      "Condition": {"StringLike": {"iam:AWSServiceName": "compute-optimizer.amazonaws.com"}}
    },
    {
      "Effect": "Allow",
      "Action": "iam:PutRolePolicy",
      "Resource": "arn:aws:iam::*:role/aws-service-role/compute-optimizer.amazonaws.com/AWSServiceRoleForComputeOptimizer"
    }
  ]
}
```

### To allow an IAM entity to create any service-linked role

Add the following statement to the permissions policy for the IAM entity that needs to create a service-linked role, or any service role that includes the needed policies. This policy attaches a policy to the role.

```
{
  "Effect": "Allow",
  "Action": "iam:CreateServiceLinkedRole",
  "Resource": "arn:aws:iam::*:role/aws-service-role/*"
}
```

## Creating a Service-Linked Role for Compute Optimizer

You don't need to manually create a service-linked role. When you opt in to the Compute Optimizer service in the AWS Management Console, the AWS CLI, or the AWS API, Compute Optimizer creates the service-linked role for you.

### **Important**

This service-linked role can appear in your account if you completed an action in another service that uses the features supported by this role. For more information, see [A New Role Appeared in My IAM Account](#).

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you opt in to the Compute Optimizer service, Compute Optimizer creates the service-linked role for you again.

## Editing a Service-Linked Role for Compute Optimizer

Compute Optimizer does not allow you to edit the `AWSServiceRoleForComputeOptimizer` service-linked role. After you create a service-linked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see [Editing a Service-Linked Role](#) in the *IAM User Guide*.

## Deleting a Service-Linked Role for Compute Optimizer

We recommend that you delete the `AWSServiceRoleForComputeOptimizer` service-linked role if you no longer need to use Compute Optimizer. That way you don't have an unused entity that is not actively monitored or maintained. However, you must opt out of Compute Optimizer before you can manually delete the service-linked role.

### **To opt out of Compute Optimizer**

For information about opting out of Compute Optimizer, see [Opting out your account \(p. 5\)](#).

### **To manually delete the service-linked role using IAM**

Use the IAM console, the AWS CLI, or the AWS API to delete the `AWSServiceRoleForComputeOptimizer` service-linked role. For more information, see [Deleting a Service-Linked Role](#) in the *IAM User Guide*.

## Supported Regions for Compute Optimizer Service-Linked Roles

Compute Optimizer supports using service-linked roles in all of the Regions where the service is available. For more information, see [AWS Regions and Endpoints](#).

## Amazon S3 bucket policy for Compute Optimizer

You can export your Compute Optimizer recommendations in a comma-separated values (.csv) file, and its metadata in a JavaScript Object Notation (.json) file, to an Amazon Simple Storage Service (Amazon S3) bucket. For more information, see [Exporting recommendations \(p. 24\)](#).

You must create the destination S3 bucket for your recommendations export before you create the export job. Compute Optimizer does not create the S3 bucket for you. The S3 bucket that you specify for your recommendations export files cannot be publicly accessible, and cannot be configured as a [Requester Pays](#) bucket.

As a best practice, create a dedicated S3 bucket for Compute Optimizer export files. For more information, see [How Do I Create an S3 Bucket?](#) in the *Amazon S3 Console User Guide*. After you create

the S3 bucket, ensure that it has the required permission policy to allow Compute Optimizer to write the export files to it. For more information, see [Specifying an existing bucket for your recommendations export](#) (p. 11).

## Using encrypted S3 buckets for your recommendations export

For the destination of your Compute Optimizer recommendations exports, you can specify S3 buckets that are encrypted with either Amazon S3-Managed Keys (SSE-S3) or Customer Master Keys (CMKs) stored in the AWS Key Management Service (AWS KMS).

You must create a symmetric CMK to use an S3 bucket with AWS KMS encryption enabled. Symmetric CMKs are the only CMKs supported by Amazon S3. For more information, see [Creating keys](#) in the *AWS KMS Developer Guide*. After you create the CMK, you must apply it to the S3 bucket that you plan to use for your recommendations export. For more information, see [How do I enable default encryption for an Amazon S3 bucket?](#) in the *Amazon S3 Console User Guide*.

Use the following procedure to grant Compute Optimizer the required permission to use your CMK to encrypt your recommendations export file when saving it to your encrypted S3 bucket.

1. Open the AWS KMS console at <https://console.aws.amazon.com/kms>.
2. To change the AWS Region, use the Region selector in the upper-right corner of the page.
3. In the left navigation menu, choose **Customer Managed Keys**.
4. Choose the name of the CMK that you chose to encrypt the export S3 bucket.
5. Choose the **Key policy** tab, then choose **Switch to policy view**.
6. Choose **Edit** to edit the key policy.
7. Copy and paste the following into the statements section of the key policy:

The statement (for the GenerateDataKey action) allows Compute Optimizer to call the AWS KMS API to obtain the data key for encrypting the recommendation files. In this way, the uploaded data format can accommodate the bucket encryption setting. Otherwise, Amazon S3 will reject the export request.

### Note

If the existing CMK already has one or more policies attached, add the statements for Compute Optimizer access to those policies. Evaluate the resulting set of permissions to be sure that they are appropriate for the users who will access the CMK.

```
{
  "Sid": "Allow use of the key to Compute Optimizer",
  "Effect": "Allow",
  "Principal": {
    "Service": "compute-optimizer.amazonaws.com"
  },
  "Action": "kms:GenerateDataKey",
  "Resource": "*"
}
```

## Specifying an existing bucket for your recommendations export

Use the following procedure to add a policy to your S3 bucket that allows Compute Optimizer to write recommendations export files to your bucket.

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Choose the bucket where you want Compute Optimizer to deliver your export files.
3. Choose **Permissions**.
4. Choose **Bucket Policy**.
5. Copy the following policy, and paste it into the **Bucket Policy Editor** text box.

Replace the placeholders in italics with the name of your bucket, the optional object prefix, and the account number of the requester of the export job. If you plan to specify an object prefix when you create your recommendations export, include it in the policy. The object prefix is an optional addition to the S3 object key that organizes your export files in your S3 bucket.

You must copy and paste this policy to include all three statements. The first statement (for the `GetBucketAcl` action) allows Compute Optimizer to get the access control list (ACL) of your bucket. The second statement (for the `GetBucketPolicyStatus` action) allows Compute Optimizer to get the policy status of your bucket, indicating whether the bucket is public. The third statement (for the `PutObject` action) gives Compute Optimizer full control to put the export file in your bucket. Your export request will fail if any of these statements is missing, or if the bucket name and optional object prefix in the policy don't match what you specify in your export request, or if the account number in the policy doesn't match the account number of the requester of the export job.

#### Note

If the existing bucket already has one or more policies attached, add the statements for Compute Optimizer access to that policy or policies. Evaluate the resulting set of permissions to be sure that they are appropriate for the users who will access the bucket.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {"Service": "compute-optimizer.amazonaws.com"},
      "Action": "s3:GetBucketAcl",
      "Resource": "arn:aws:s3:::myBucketName"
    },
    {
      "Effect": "Allow",
      "Principal": {"Service": "compute-optimizer.amazonaws.com"},
      "Action": "s3:GetBucketPolicyStatus",
      "Resource": "arn:aws:s3:::myBucketName"
    },
    {
      "Effect": "Allow",
      "Principal": {"Service": "compute-optimizer.amazonaws.com"},
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::myBucketName/optional prefix/compute-optimizer/myAccountID/*",
      "Condition": {"StringEquals": {"s3:x-amz-acl": "bucket-owner-full-control"}}
    }
  ]
}
```

If you don't want to specify an object prefix, use the following policy instead.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {"Service": "compute-optimizer.amazonaws.com"},
```



```
        "Action": "s3:GetBucketAcl",
        "Resource": "arn:aws:s3:::myBucketName"
    },
    {
        "Effect": "Allow",
        "Principal": {"Service": "compute-optimizer.amazonaws.com"},
        "Action": "s3:GetBucketPolicyStatus",
        "Resource": "arn:aws:s3:::myBucketName"
    },
    {
        "Effect": "Allow",
        "Principal": {"Service": "compute-optimizer.amazonaws.com"},
        "Action": "s3:PutObject",
        "Resource": "arn:aws:s3:::myBucketName/compute-optimizer/myAccountID/*",
        "Condition": {"StringEquals": {"s3:x-amz-acl": "bucket-owner-full-
control"}}
    }
]
}
```

## Additional resources

For more information about S3 buckets and policies, see the [Amazon Simple Storage Service Console User Guide](#).

# Metrics analyzed by AWS Compute Optimizer

After you [opt in \(p. 4\)](#), AWS Compute Optimizer begins analyzing the specifications (vCPUs, memory, storage, and so on) and the utilization metrics of your running resources from a period of the last 14 days. Compute Optimizer requires at least 30 *consecutive* hours of metrics data from your resource to generate recommendations. After the analysis is completed, which could take up to 12 hours, Compute Optimizer presents its findings on the dashboard page. For more information, see [Viewing the AWS Compute Optimizer dashboard \(p. 16\)](#).

Compute Optimizer analyzes the following CloudWatch metrics of your EC2 instances by default:

- **CPU utilization** - The percentage of allocated EC2 compute units that are in use on the instance. This metric identifies the processing power required to run an application on an instance.
- **Memory utilization** - The amount of memory that has been used in some way during the sample period. This metric identifies the memory required to run an application on an instance. 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 CloudWatch Agent \(p. 14\)](#).
- **Network in** - The number of bytes received on all network interfaces by the instance. This metric identifies the volume of incoming network traffic to an instance.
- **Network out** - The number of bytes sent out on all network interfaces by the instance. This metric identifies the volume of outgoing network traffic from an instance.
- **Amazon EBS volume read/write throughput and bandwidth** - The read/write operations per second, and the read/write bytes per second (displayed as KiB/seconds in the console) of all EBS volumes attached to the instance.

For more information about these metrics, see [List the Available CloudWatch Metrics for Your Instances](#) in the *Amazon Elastic Compute Cloud User Guide*.

## Enabling memory utilization with the CloudWatch Agent

Install the CloudWatch agent on your instances to have Compute Optimizer analyze the memory utilization of your instances. Enabling Compute Optimizer to analyze memory utilization data for your instances provides an additional measurement of data that further improves the recommendations provided by the service. For more information about installing the CloudWatch agent, see [Collecting Metrics and Logs from Amazon EC2 Instances and On-Premises Servers with the CloudWatch Agent](#) in the *Amazon CloudWatch 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. Additionally, the namespace must contain the `InstanceId` dimension. Compute Optimizer will not be able to collect memory utilization data for your instance if the `InstanceId` dimension is missing, or if you overwrite it with your own custom dimension name. Namespaces and dimensions are defined in the CloudWatch Agent configuration file. For more information, see [Create the CloudWatch Agent Configuration File](#) in the *Amazon CloudWatch User Guide*.

## Enabling detailed monitoring in CloudWatch

Enable detailed monitoring for your resources to allow your instance resources to be monitored at 1-minute increments, rather than the default 5-minute increments. Detailed monitoring provides a higher level of definition for the utilization data that is available to Compute Optimizer. For more information, see [Enable or Disable Detailed Monitoring for Your Instances](#) in the *Amazon Elastic Compute Cloud User Guide*, or [Configure Monitoring for Auto Scaling Instances](#) in the *Amazon EC2 Auto Scaling User Guide*. Additional charges might apply to enable detailed monitoring. For more information, see the [Amazon CloudWatch pricing page](#).

# Viewing the AWS Compute Optimizer dashboard

The dashboard in the AWS Compute Optimizer console displays an overview of optimization findings for your AWS resources. Your resources are organized into findings classifications, and are separated into service-specific sections. For example, the dashboard shows the findings for your EC2 instances and your Auto Scaling groups separately. Use the dashboard to identify the optimization opportunities under each service that is supported by Compute Optimizer.

The findings and recommendations are refreshed daily. They're generated by analyzing your resource's specifications and utilization metrics over a period of the last 14 days. For more information, see [Metrics analyzed by AWS Compute Optimizer \(p. 14\)](#). Keep in mind that findings and recommendations could take up to 12 hours to be generated, and sufficient metric data must be accumulated. For more information, see [Metric requirements \(p. 3\)](#).

## Findings classifications

The findings classifications in the Compute Optimizer console provide an at-a-glance view of how your resources performed during the analyzed period. A percentage, and a count of your resources that are in each classification, are shown on the dashboard. The classifications differ based on the AWS service of the resource. The following classifications apply to services and resources supported by Compute Optimizer.

### EC2 instances

- **Under-provisioned** - An EC2 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 EC2 instances might lead to poor application performance.
- **Over-provisioned** - An EC2 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 when no specification is under-provisioned. Over-provisioned EC2 instances might lead to unnecessary infrastructure cost.
- **Optimized** - An EC2 instance is considered optimized when all specifications of your instance, such as CPU, memory, and network, meet the performance requirements of your workload, and the instance is not over-provisioned. An optimized EC2 instance runs your workloads with optimal performance and infrastructure cost. For optimized resources, Compute Optimizer might sometimes recommend a new generation instance type.

### Auto Scaling groups

- **Not optimized** - An Auto Scaling group is considered not optimized when Compute Optimizer has identified a recommendation that can provide better performance or cost 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 sometimes recommend a new generation instance type.

## Viewing the dashboard

Use the following procedure to view the dashboard and the optimization findings for your resources.

### To view the dashboard

1. Open the Compute Optimizer console at <https://console.aws.amazon.com/compute-optimizer/>.
2. Choose **Dashboard** in the navigation pane.

By default, the dashboard displays an overview of optimization findings for AWS resources across all AWS Regions in the account that you are currently signed in to.

3. You can perform the following actions on the dashboard:
  - To view the optimization findings for resources in another account, choose **Account**, and then select a different account ID.

If you are signed in to a management account of an organization, and you opted in only your account (rather than opting in all member accounts within the organization), you will get an error if you try to view the optimization findings for resources in another account (within your organization that has not yet opted in to Compute Optimizer).
  - To filter findings on the dashboard to one or more AWS Regions, enter the name of the Region in the **Filter by one or more Regions** text box, or choose one or more Regions in the drop-down list that appears.
  - To clear the selected filters, choose **Clear filters** next to the filter.
  - To view optimization recommendations, choose the **View recommendations** link for one of the resource types displayed, or choose the number of resources listed next to a findings classification to view the resources for that classification. For more information, see [Viewing resource recommendations \(p. 18\)](#).

# Viewing resource recommendations

The recommendations page in the AWS Compute Optimizer console lists your running resources, their [findings classifications](#) (p. 18), instance type, and hourly prices for the selected [purchasing option](#) (p. 18). The top recommendation from Compute Optimizer is listed next to each of your resources, and it includes the recommended instance type and the hourly price for the selected purchasing option. Use the recommendations page to compare your running resources with recommended resources, which can help you to decide if you should up-size or down-size your resources.

The recommendations are refreshed daily. They're generated by analyzing your resource's specifications and utilization metrics over a period of the last 14 days. For more information, see [Metrics analyzed by AWS Compute Optimizer](#) (p. 14). Keep in mind that recommendations could take up to 12 hours to be generated, and sufficient metric data must be accumulated. For more information, see [Metric requirements](#) (p. 3).

## Supported resources and requirements

Compute Optimizer generates recommendations for Amazon Elastic Compute Cloud (Amazon EC2) instances and Amazon EC2 Auto Scaling groups that meet a specific set of requirements, and that have accumulated sufficient metric data. For more information, see [Supported resources and requirements](#) (p. 3).

## Findings classifications

The findings classifications on the recommendations page provide a summary of how each of your resources performed during the analyzed period, such as whether they were over-provisioned or under-provisioned. The classifications differ based on the AWS service of the resource. For more information, see [Findings classifications](#) (p. 16).

## Prices and purchasing options

In the recommendations page, you can choose to view the hourly prices for your running instances and the recommended instance types under different Amazon EC2 purchasing options. Use the pricing information to get an idea of the price difference between your running resources and the recommended resources. It also shows the price difference between your running instance's current purchasing option, and the purchasing option that you select on the recommendations page.

### **Important**

The prices listed on the recommendations page might not reflect the actual prices you pay for your resources. For more information about finding the actual price of your running resources, see [Amazon EC2 Usage Reports](#) in the *Amazon Elastic Compute Cloud User Guide*.

The following purchasing options can be selected on the recommendations page:

- **On-Demand Instances** - An On-Demand Instance is an instance that you use on demand. You have full control over its lifecycle—you decide when to launch, stop, hibernate, start, reboot, and terminate it. No longer-term commitments or upfront payments are needed. For more information about On-Demand Instances, see [On-Demand Instances](#) in the *Amazon Elastic Compute Cloud User Guide*. For more information about pricing, see [Amazon EC2 On-Demand Instance Pricing](#).

- **Reserved Instances (one-year or three-year commitment, no upfront)** - Reserved Instances provide you with significant savings on your Amazon EC2 costs compared to On-Demand Instance pricing. Reserved Instances are not physical instances, but rather a billing discount applied to the use of On-Demand Instances in your account. For more information about Reserved Instances, see [Reserved Instances](#) in the *Amazon Elastic Compute Cloud User Guide*. For more information about pricing, see [Amazon EC2 Reserved Instance Pricing](#).

For more information about purchasing options, see [Instance Purchasing Options](#) in the *Amazon Elastic Compute Cloud User Guide*.

## Estimated monthly savings

The EC2 instances recommendations page also lists the **On-Demand hours**, **Reserved Instance (RI) hours**, and **Savings Plans hours** for each of your running instances. These data columns represent the subset of the total running hours that each instance was On-Demand, was covered by an AWS reservation, and was covered by Savings Plans over the look-back period (up to the last 14 days). That data is used to calculate an **Estimated monthly savings** for each of your recommendations, which is also displayed in the recommendations page.

To calculate the estimated monthly savings, we first examine the instance running in the look-back period to identify if it was partially or fully covered by an RI or running On-Demand. Another factor is whether the RI is size-flexible. The cost to run the instance is calculated based on the On-Demand hours and the rate of the instance type.

For each recommendation, we calculate the cost to operate a new instance. We assume that a size-flexible RI will cover the new instance in the same way as the previous instance. Savings are calculated based on the number of On-Demand running hours and the difference in On-Demand rates. If the RI isn't size-flexible, the savings calculation is based on if the instance hours during the look-back period are operated as On-Demand.

Recommendations are provided regardless of estimated monthly savings. Even if the estimated savings is less than zero, you might be able to convert your existing Convertible RIs to cover the recommended instance types.

### Note

Recommendations don't capture second-order effects of rightsizing, such as the resulting RI hour's availability and how they will apply to other instances. Potential savings based on reallocation of the RI hours aren't included in the calculation.

## Viewing recommendations

Use the following procedure to view recommendations for your resources.

### To view recommendations

1. Open the Compute Optimizer console at <https://console.aws.amazon.com/compute-optimizer/>.
2. Choose one of the recommendation page options listed in the navigation pane, such as **EC2 instances** or **Auto Scaling groups**.

The recommendations page lists the specifications and findings classifications of your running resources, along with the specifications of the recommended resources. The resources listed are from the AWS Region that is currently selected, in the selected account.

3. You can perform the following actions on the recommendations page:

- To filter recommendations to one or more AWS Regions, enter the name of the Region in the **Filter by one or more Regions** text box, or choose one or more Regions in the drop-down list that appears.
- To view recommendations for resources in another account, choose **Account**, and then select a different account ID.

This option is available only if you are signed in to a management account of an organization, and you opted in all member accounts within the organization.

- To clear the selected filters, choose **Clear filters** next to the filter.
- To change the purchasing option that is displayed, choose **Settings** (the gear icon), then choose **On-Demand Instances**, **Reserved Instances, standard 1-year no upfront**, or **Reserved Instances, standard 3-year no upfront**.
- To view the resource details, such as additional recommendations and a comparison of utilization metrics, choose the findings classification listed next to the desired resource. For more information, see [Viewing resource details \(p. 21\)](#).



# Viewing resource details

The resource details page in the AWS Compute Optimizer console lists multiple optimization recommendations for a specific AWS resource running in your account. It lists the specifications (for example, vCPUs, memory, or storage) of your running instance, along with the specifications and performance risks of the recommended instance types. The resource details page also displays graphs that show the following:

- Utilization metrics of your running instance during the analyzed period (up to 14 days).
- A CPU and memory utilization comparison with the recommended resource that is selected in the list.
- The read/write performance of the Amazon EBS volumes attached to the instance.

You can use the resource details page to see how your running resource performed during the analyzed period, and to view all of the recommendations generated by Compute Optimizer for that resource.

The recommendations are refreshed daily. They're generated by analyzing your resource's specifications and utilization metrics over a period of the last 14 days. For more information, see [Metrics analyzed by AWS Compute Optimizer \(p. 14\)](#). Keep in mind that recommendations could take up to 12 hours to be generated, and sufficient metric data must be accumulated. For more information, see [Metric requirements \(p. 3\)](#).

## Performance risk

Each of the recommendations listed on the resource details page includes a performance risk classification (high, medium, low, or very low). The risk classification indicates the likelihood that the recommended resource does not meet the performance requirements of your workload. When comparing recommendations, and determining which one best suits the needs of your workload, carefully weigh the possible performance risks against the benefits. Decide whether to optimize for performance improvement, for cost reduction, or for a combination of these two. For more information, see [Changing the Instance Type](#) in the *Amazon Elastic Compute Cloud User Guide*.

## Prices and purchasing options on the resource details page

The resource details page lists the estimated price of your running resource and the price of the recommended resources under different pricing options. For more information, see [Prices and purchasing options \(p. 18\)](#).

## Utilization graphs

The resource details page displays utilization metric graphs for your running resource. The graphs display data for the analyzed period (up to 14 days). You can change the graphs to display data for the last 24 hours, three days, one week, or two weeks. You can also change the statistic of the graphs between average and maximum.

The following utilization graphs are displayed on the resource details page:

Graph name	Description
CPU utilization	The percentage of allocated EC2 compute units used by the instance.
EBS read bandwidth	The read kibibytes (KiB) per second from all EBS volumes attached to the instance.
EBS read throughput	The completed read operations per second from all EBS volumes attached to the instance.  For Xen instances, data is reported only when there is read activity on the volume.
EBS write bandwidth	The written kibibytes (KiB) per second to all EBS volumes attached to the instance.
EBS write throughput	The completed write operations per second to all EBS volumes attached to the instance.  For Xen instances, data is reported only when there is write activity on the volume.
Memory utilization	The percentage of memory allocated by applications and the operating system as used.  The memory utilization graph is populated only for your resources that have the unified CloudWatch agent installed on them. For more information, see <a href="#">Collecting Metrics and Logs from Amazon EC2 Instances and On-Premises Servers with the CloudWatch Agent</a> in the <i>Amazon CloudWatch User Guide</i> .
Network in	The number of bytes received on all network interfaces by the instance.
Network out	The number of bytes sent out on all network interfaces by the instance.

The CPU utilization and memory utilization graphs include a comparison of your CPU and memory utilization data against that of the selected recommended resource. The comparison shows you what your CPU and memory utilization would have been if you had used the recommended resource during the analyzed period. This comparison can help you to identify if the recommended resource is within your workload's performance threshold.

## Viewing resource details

Use the following procedure to view the details of your resource.

### To view resource details

1. Open the Compute Optimizer console at <https://console.aws.amazon.com/compute-optimizer/>.
2. Choose one of the recommendation options listed in the navigation pane, such as **EC2 instances** or **Auto Scaling groups**.

3. Choose the findings classification listed next to a resource.

The resource details page lists multiple optimization recommendations for the resource that you chose. It lists the specifications of your running resource, the specifications and performance risks of the recommended resources, and utilization graphs.

4. You can perform the following actions on the resource details page:

- To view the utilization comparison between your resource and a recommended resource, choose a recommended option.

The utilization metric graphs for your running resource are displayed at the bottom of the page. The solid blue line is the utilization of your running instance. The dotted orange line, displayed in the CPU utilization and memory utilization graphs, is the projected utilization of the selected recommended resource if you had used that resource during the analyzed period.

- To change the time range of the graphs, choose **Time Range**, and then choose **Last 24 hours**, **Last 3 days**, **Last week**, or **Last 2 weeks**.

Choosing a shorter time range displays the data points at a higher granularity, which provides a higher level of detail.

- To change the statistic value of the graphs, choose **Statistics**, and then choose **Average** or **Maximum**.

You can use this option to determine the typical resource utilization of your workload over time. To view the highest value observed during the specified period, change the selection to **Maximum**. This allows you to determine the peak resource usage of your workload over time.

- To change the purchasing option that is displayed, choose **Settings** (the gear icon), then choose **On-Demand Instances**, **Reserved Instances, standard 1-year no upfront**, or **Reserved Instances, standard 3-year no upfront**.

# Exporting recommendations

You can export your recommendations to record them over time, and share the data with others. 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.

## Contents

- [Restrictions](#) (p. 24)
- [Amazon S3 bucket permission requirements](#) (p. 24)
- [Exporting your recommendations](#) (p. 24)
- [Viewing your export jobs](#) (p. 26)
- [Exported files](#) (p. 26)

## Restrictions

The following restrictions apply when exporting recommendations:

- You can have only one recommendations export job in progress for each resource type, and for each AWS Region. Confirm that all previous export jobs are completed before creating a new export job. For more information about viewing your export jobs, including those that are in progress, see [Viewing your export jobs](#) (p. 26).
- Recommendations for each resource type, and in each AWS Region are exported in separate .csv files. You cannot export recommendations from multiple resource types and Regions into a single file.
- Large export jobs can take a while to complete. Consider limiting the recommendation columns that you include in your export job. Additionally, if your account is the management account of an organization, consider limiting the number of member accounts to include in your export job.

## Amazon S3 bucket permission requirements

You must create the destination S3 bucket for your recommendations export before you create the export job. Compute Optimizer does not create the S3 bucket for you. The S3 bucket that you specify for your recommendations export files cannot be publicly accessible, and cannot be configured as a [Requester Pays](#) bucket. After you create the S3 bucket, ensure that it has the required permission policy to allow Compute Optimizer to write the export files to it. If you plan to specify an object prefix when you create your recommendations 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](#) (p. 10).

## Exporting your recommendations

Use the following procedure to export your recommendations.

### To export your recommendations

1. Open the Compute Optimizer console at <https://console.aws.amazon.com/compute-optimizer/>.
2. Choose a resource type in the navigation pane. For example, choose **EC2 instances**, or **Auto Scaling groups**.
3. On the **Recommendations** page, choose the **Action** dropdown menu, and choose **Export Recommendations**.

4. On the **Export Recommendations** page, under **Export destination settings**, specify the following:
  - a. For **Region**, specify an AWS Region for your export.
  - b. For **Destination S3 bucket name**, specify the name of an existing S3 bucket in the specified Region.
  - c. (Optional) Choose **Add Region** to export the recommendations for another AWS Region.
  - d. (Optional) Choose **Remove** next to a specified Region and S3 bucket name to remove the destination from the export job.
  - e. (Optional) For **Object prefix**, specify a prefix to use in the destination S3 bucket for all of the export files. The prefix is an optional addition to the S3 object key that organizes your export files in your S3 bucket. For example, you can specify a date prefix (e.g., `2020/april`), a resource type prefix (e.g., `ec2-instances`), or a combination of both (e.g., `2020/april/ec2-instances`).
5. Under **Export filters**, specify the following:
  - a. For **Resource type**, choose the resource type to include in your recommendations export.
  - b. For **Accounts**, choose if you want to include recommendations for all member accounts of the organization. This option is available only if your account is the management account of an organization.
6. Under **Columns to include**, choose the recommendations data to include in your recommendations export. For more information about the columns to include, see [Exported files \(p. 26\)](#).
7. Choose **Export** after you confirm that the export job is configured correctly, or choose **Cancel** to return to the **Recommendations** page without creating the export job. The export job configuration is deleted if you cancel it.

**Note**

If you export recommendations for multiple AWS Regions at one time, they are treated as separate export jobs and Compute Optimizer will try to start all of them at once. The **Export Recommendations** page will display an error if an export job fails to start. Export jobs that successfully start will continue to process, but you must resolve the errors for the failed jobs before trying to start them again.

Your recommendations export job will take a while to complete. Check the status of your export jobs by viewing the **Exports** page. For more information, see [Viewing your export jobs \(p. 26\)](#). Your recommendations export file, and its associated metadata file, are saved to the specified S3 bucket when the export job is completed. The following are examples of the full Amazon S3 object key for the export file and its associated metadata file. The account ID in the object keys is the account of the requester of the export job. For more information, see [Exported files \(p. 26\)](#).

```
s3://<BucketName>/<OptionalPrefix>/compute-optimizer/<AccountId>/<AWSRegion>-<CreatedTimestamp>-<UniqueJobID>.csv
```

```
s3://<BucketName>/<OptionalPrefix>/compute-optimizer/<AccountId>/<AWSRegion>-<CreatedTimestamp>-<UniqueJobID>-metadata.json
```

**Example:**

```
s3://compute-optimizer-exports/ec2-instance-recommendations/compute-optimizer/111122223333/us-west-2-2020-03-03T133027-3e496c549301c8a4dfcsdX.csv
```

```
s3://compute-optimizer-exports/ec2-instance-recommendations/compute-optimizer/111122223333/us-west-2-2020-03-03T133027-3e496c549301c8a4dfcsdX-metadata.json
```

## Viewing your export jobs

Use the following procedure to view the export jobs created in the last seven days.

### To view your export jobs

1. Open the Compute Optimizer console at <https://console.aws.amazon.com/compute-optimizer/>.
2. Choose **Exports** in the navigation pane.

The **Exports** page displays the recommendation export jobs that were created in the last seven days.

3. You can perform the following actions on the page:
  - Choose the export destination link for a completed job to access the destination S3 bucket. The export destination displays only for successful export jobs. Export jobs that are in progress, or that have failed display a dash.
  - Scroll right to view the failure reason for failed export jobs. Use the failure reason to determine why your export job was not completed.

## Exported files

Recommendations are exported in a .csv file, and its metadata in a .json file, to the Amazon S3 bucket that you specified when you created the export job.

### Recommendations file

The recommendations file includes the recommendation data for the recommendation columns that you choose to include when you create the export job. The following table lists all of the recommendation columns that can be included in the export file.

In the table, *Column name* represents the fields that are exported in the .csv file, and *Description* describes the data of each field and the name of the column as displayed in the Compute Optimizer console. The **EC2 instances** and **Auto Scaling groups** columns specify if the data applies to each of those resource types. The recommendation data columns in the .csv file are numbered when multiple recommendations are generated for each resource. Ranked recommendation columns, in which *<rank>* is replaced with a ranking, correspond to each other. For example, recommendationOptions\_1\_memory, recommendationOptions\_1\_network, and recommendationOptions\_1\_vcpus correspond to each other, and are for the same recommendation.

Column name	Description	EC2 instances	Auto Scaling groups
accountId	The AWS account ID in which the source resource was created. This column is listed as <b>Account ID</b> in the Compute Optimizer console.	Yes	Yes
autoScalingGroupArn	The Amazon Resource Name (ARN) of the source Auto Scaling group. This column is listed as <b>Auto Scaling group ARN</b> in the	No	Yes

Column name	Description	EC2 instances	Auto Scaling groups
	Compute Optimizer console.		
autoScalingGroupName	The name of the Auto Scaling group. This column is listed as <b>Auto Scaling group name</b> in the Compute Optimizer console.	No	Yes
currentConfiguration_desiredCapacity	The capacity of the source Auto Scaling group. This column is listed as <b>Current desired capacity</b> in the Compute Optimizer console.	No	Yes
currentConfiguration_instanceType	The instance type of the source Auto Scaling group. This column is listed as <b>Current instance type</b> in the Compute Optimizer console.	No	Yes
currentConfiguration_maximumSize	The maximum size of the source Auto Scaling group. This column is listed as <b>Current maximum size</b> in the Compute Optimizer console.	No	Yes
currentConfiguration_minimumSize	The minimum size of the source Auto Scaling group. This column is listed as <b>Current minimum size</b> in the Compute Optimizer console.	No	Yes
currentInstanceType	The instance type of the source instance. This column is listed as <b>Current instance type</b> in the Compute Optimizer console.	Yes	No

Column name	Description	EC2 instances	Auto Scaling groups
current_memory	The memory of the source instance, including for instances in an Auto Scaling group. This column is listed as <b>Current memory</b> in the Compute Optimizer console.	Yes	Yes
current_network	The network performance, or rate of data transfer, of the source instance, including for instances in an Auto Scaling group. This column is listed as <b>Current network</b> in the Compute Optimizer console.	Yes	Yes
current_onDemandPrice	The On-Demand price of the source instance, including for instances in an Auto Scaling group. The price listed might not reflect the actual price you pay for the instance. This column is listed as <b>Current On-Demand price</b> in the Compute Optimizer console.	Yes	Yes
current_standardOneYearNoUpfrontPrice	The price of the source instance, including for instances in an Auto Scaling group. The price listed might not reflect the actual price you pay for the instance. This column is listed as <b>Current 1-year RI price</b> in the Compute Optimizer console.	Yes	Yes



Column name	Description	EC2 instances	Auto Scaling groups
current_standardThreeYearUpfrontPrice	The upfront price of a standard 3-year no upfront price of the source instance, including for instances in an Auto Scaling group. The price listed might not reflect the actual price you pay for the instance. This column is listed as <b>Current 3-year RI price</b> in the Compute Optimizer console.	Yes	Yes
current_storage	The local storage volume of the source instance, including for instances in an Auto Scaling group. This column is listed as <b>Current storage</b> in the Compute Optimizer console.	Yes	Yes
current_vcpus	The number of vCPUs of the source instance, including for instances in an Auto Scaling group. This column is listed as <b>Current vCPUs</b> in the Compute Optimizer console.	Yes	Yes
finding	The finding classification for the source instance, or for the source Auto Scaling group. Instances can be classified as under-provisioned, optimized, or over-provisioned. Auto Scaling groups can be classified as not optimized, or optimized. This column is listed as <b>Finding</b> in the Compute Optimizer console.	Yes	Yes

Column name	Description	EC2 instances	Auto Scaling groups
instanceArn	The Amazon Resource Name (ARN) of the source instance. This column is listed as <b>Instance ARN</b> in the Compute Optimizer console.	Yes	No
instanceName	The name of the source instance. This column is listed as <b>Instance name</b> in the Compute Optimizer console.	Yes	No
lastRefreshTimestamp	The timestamp of when the resource recommendation was last refreshed. This column is listed as <b>Last refresh timestamp</b> in the Compute Optimizer console.	Yes	Yes
lookBackPeriodInDays	The number of preceding days for which Compute Optimizer analyzed metric data from the source resource to generate the recommendation. This column is listed as <b>Lookback period in days</b> in the Compute Optimizer console.	Yes	Yes
recommendationOptions_desiredCapacity	The desired capacity of the Auto Scaling group recommendation. This column is listed as <b>Recommendation options desired capacity</b> in the Compute Optimizer console.	Yes	Yes
recommendationOptions_instanceType	The instance type of the Auto Scaling group recommendation. This column is listed as <b>Recommendation options instance type</b> in the Compute Optimizer console.	Yes	Yes

Column name	Description	EC2 instances	Auto Scaling groups
recommendationOptions_maxConfigurationSize	The maximum size of the Auto Scaling group recommendation. This column is listed as <b>Recommendation options maximum size</b> in the Compute Optimizer console.	Yes	Yes
recommendationOptions_minConfigurationSize	The minimum size of the Auto Scaling group recommendation. This column is listed as <b>Recommendation options minimum size</b> in the Compute Optimizer console.	Yes	Yes
recommendationOptions_instanceType	The instance type of the instance recommendation. This column is listed as <b>Recommendation options Instance type</b> in the Compute Optimizer console.	Yes	No
recommendationOptions_memory	The memory of the instance recommendation, including for instances in an Auto Scaling group. This column is listed as <b>Recommendation options memory</b> in the Compute Optimizer console.	Yes	Yes
recommendationOptions_network	The network performance, or rate of data transfer, of the instance recommendation, including for instances in an Auto Scaling group. This column is listed as <b>Recommendation options network</b> in the Compute Optimizer console.	Yes	Yes

Column name	Description	EC2 instances	Auto Scaling groups
recommendationOptions_TheOnDemandPrice	The price of the instance recommendation, including for instances in an Auto Scaling group. This column is listed as <b>Recommendation options On-Demand price</b> in the Compute Optimizer console.	Yes	Yes
recommendationOptions_ThePerformanceRisk	The performance risk of the instance recommendation, including for instances in an Auto Scaling group. This column is listed as <b>Recommendation options performance risk</b> in the Compute Optimizer console.	Yes	Yes
recommendationOptions_TheProjectedUtilizationMetrics_CPU_MAXIMUM	The projected maximum CPU utilization metric of the instance recommendation, including for instances in an Auto Scaling group. This value defines the maximum CPU utilization if you had used the recommended resource during the look-back period. This column is listed as <b>Recommendation options projected utilization metrics CPU maximum</b> in the Compute Optimizer console.	Yes	Yes

Column name	Description	EC2 instances	Auto Scaling groups
recommendationOptions_ <del>The projected</del> UtilizationMetrics_MEMORY_MAXIMUM	<p><del>The projected</del> maximum memory utilization metric of the instance recommendation, including for instances in an Auto Scaling group. This value defines the maximum memory utilization if you had used the recommended resource during the look-back period. This column is listed as <b>Recommendation options projected utilization metrics memory maximum</b> in the Compute Optimizer console.</p>	Yes	Yes
recommendationOptions_ <del>The Reserved</del> OneYearMarkUpfrontReservedPrice	<p><del>The Reserved</del> standard 1-year no upfront price for the instance recommendation, including for instances in an Auto Scaling group. This column is listed as <b>Recommended options 1-year RI price</b> in the Compute Optimizer console.</p>	Yes	Yes
recommendationOptions_ <del>The Reserved</del> ThreeYearMarkUpfrontReservedPrice	<p><del>The Reserved</del> standard 3-year no upfront price for the instance recommendation, including for instances in an Auto Scaling group. This column is listed as <b>Recommended options 3-year RI price</b> in the Compute Optimizer console.</p>	Yes	Yes

Column name	Description	EC2 instances	Auto Scaling groups
recommendationOptions_storage	The local storage volume of the instance recommendation, including for instances in an Auto Scaling group. This column is listed as <b>Recommendation options storage</b> in the Compute Optimizer console.	Yes	Yes
recommendationOptions_vcpus	The vCPUs of the instance recommendation, including for instances in an Auto Scaling group. This column is listed as <b>Recommendation options vCPUs</b> in the Compute Optimizer console.	Yes	Yes
recommendationsSources_AmazonResourceName	The Amazon Resource Name (ARN) of the source resource. This column is listed as <b>Recommendation source ARN</b> in the Compute Optimizer console.	Yes	No
recommendationsSources_sourceType	The source recommendation resource type. This column is listed as <b>Recommendation source type</b> in the Compute Optimizer console.	Yes	No
utilizationMetrics_CPU_Maximum	The maximum CPU utilization metric of the source instance, including for instances in an Auto Scaling group. This column is listed as <b>Utilization metrics CPU maximum</b> in the Compute Optimizer console.	Yes	Yes

Column name	Description	EC2 instances	Auto Scaling groups
utilizationMetrics_MEMORY_MAXIMUM	The <b>MAXIMUM</b> memory utilization metric of the source instance, including for instances in an Auto Scaling group. This column is listed as <b>Utilization metrics memory maximum</b> in the Compute Optimizer console.	Yes	Yes
UtilizationMetricsEBSReadOpsPerSecondMaximum	The <b>Maximum</b> of read operations per second for volumes attached to an instance. This column is listed as <b>EBS read (operations/second) maximum</b> in the Compute Optimizer console.	Yes	Yes
UtilizationMetricsEBSWriteOpsPerSecondMaximum	The <b>Maximum</b> of write operations per second for volumes attached to an instance. This column is listed as <b>EBS write (operations/second) maximum</b> in the Compute Optimizer console.	Yes	Yes
UtilizationMetricsEBSReadBytesPerSecondMaximum	The <b>Maximum</b> read per second for volumes attached to an instance. This column is listed as <b>EBS read (bytes/second) maximum</b> in the Compute Optimizer console.	Yes	Yes
UtilizationMetricsEBSWriteBytesPerSecondMaximum	The <b>Maximum</b> bytes written per second for volumes attached to an instance. This column is listed as <b>EBS write (bytes/second) maximum</b> in the Compute Optimizer console.	Yes	Yes

Column name	Description	EC2 instances	Auto Scaling groups
recommendations_count	The number of recommendation included. This column is not listed in the Compute Optimizer console.	Yes	Yes
errorCode	The error code for when a recommendation was not generated for a resource. This column is not listed in the Compute Optimizer console.	Yes	Yes
errorMessage	The error message that corresponds to the error code. This column is not listed in the Compute Optimizer console.	Yes	Yes

## Metadata file

A metadata .json file is output with every export job. It includes the schema information for the associated recommendations file, such as the dialect of the data, column definitions, column descriptions, and more. The file is meant to help parse the export file, and describe its contents. The metadata file is saved in the same S3 bucket and prefix that you specified for the export file.

The metadata file includes the following properties for each exported column, or field:

- **Name** - The export field recommendations column name.
- **Titles** - The user-friendly recommendations column name.
- **Datatype** - The type of data for the column.
- **Null** - The string to expect if the column is null.
- **Required** - Indicates if the column data is required.

The following is an example of the information included in the metadata file.

```
{
  "@context": [
    "http://www.w3.org/ns/csvw"
  ],
  "url": "us-east-1-2020-05-18T001229Z-f264881a-bfb3-4676-9b14-8d1243599ebb.csv",
  "dc:title": "EC2 Instance Recommendations",
  "dialect": {
    "encoding": "utf-8",
    "lineTerminators": [
      "\n"
    ],
    "doubleQuote": true,
    "skipRows": 0,
    "header": true,
    "headerRowCount": 1,
    "delimiter": ",",
  }
}
```



```
"skipColumns": 0,
"skipBlankRows": false,
"trim": false
},
"dc:modified": {
"@value": "2020-05-20",
"@type": "xsd:date"
},
"tableSchema": {
"columns": [
{
"name": "accountId",
"titles": "Account ID",
"datatype": "string",
>null": "",
"required": false
},
{
"name": "instanceArn",
"titles": "Instance Arn",
"datatype": "string",
>null": "",
"required": false
},
{
"name": "utilizationMetrics_CPU_MAXIMUM",
"titles": "Cpu Maximum Utilization Metrics",
"datatype": "double",
>null": "",
"required": false
},
{
"name": "recommendations_count",
"titles": "Number of recommendations",
"datatype": "integer",
"required": true
},
{
"name": "recommendationOptions_1_instanceType",
"titles": "Recommendation 1 Instance Type",
"datatype": "integer",
>null": "",
"required": false
},
{
"name": "lastRefreshTimestamp.UTC",
"titles": "Last Resfreshed Timestamp UTC",
"datatype": "datetime",
"format": "yyyy-MM-dd HH:mm:ss",
>null": "",
"required": false
},
{
"name": "errorCode",
"titles": "Error Code",
"datatype": "string",
"required": true
},
{
"name": "errorMessage",
"titles": "Error Message",
"datatype": "string",
"required": true
}
]
}
```

```
}
```

# Document history for AWS Compute Optimizer

The following table describes the documentation for this release of AWS Compute Optimizer.

- **API version:** 2019-11-30
- **Latest documentation update:** October 7, 2020

The following table describes the documentation for this release of Compute Optimizer.

update-history-change	update-history-description	update-history-date
<a href="#">Amazon EBS metrics and encrypted S3 buckets (p. 39)</a>	Compute Optimizer now analyzes the read/write operations per second, and the read/write bytes per second of Amazon Elastic Block Store (Amazon EBS) volumes attached to an instance. The data is used to generate recommendations. You can also view EBS read/write bandwidth (operations/second), and EBS read/write throughput (KiB/second) graphs in the Compute Optimizer console. For more information, see <a href="#">Viewing resource details</a> . You can also now export recommendations to encrypted Amazon S3 buckets. For more information, see <a href="#">Exporting recommendations</a> .	October 7, 2020
<a href="#">Recommendations export (p. 39)</a>	Recommendations can be exported to Amazon Simple Storage Service (Amazon S3). For more information, see <a href="#">Exporting recommendations</a> .	June 10, 2020
<a href="#">Self-service opt out (p. 39)</a>	Self-service opt out is now supported by using the AWS Command Line Interface. For more information, see <a href="#">Opting out your account</a> .	April 6, 2020
<a href="#">Service release (p. 39)</a>	Compute Optimizer released.	December 2, 2019