
Amazon CloudWatch

User Guide



Amazon CloudWatch: User Guide

Copyright © 2016 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

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

Table of Contents

What is Amazon CloudWatch?	1
Accessing CloudWatch	1
Related AWS Services	1
How CloudWatch Works	2
Concepts	3
Namespaces	3
Metrics	3
Dimensions	4
Statistics	5
Percentiles	7
Alarms	7
Limits	7
Resources	8
Getting Set Up	10
Sign Up for Amazon Web Services (AWS)	10
Sign in to the Amazon CloudWatch Console	10
Set Up the AWS CLI	11
Getting Started	12
Scenario: Monitor Estimated Charges	12
Step 1: Enable Billing Alerts	13
Step 2: Create a Billing Alarm	13
Step 3: Check the Alarm Status	14
Step 4: Edit a Billing Alarm	15
Step 5: Delete a Billing Alarm	15
Scenario: Publish Metrics	15
Step 1: Define the Data Configuration	16
Step 2: Add Metrics to CloudWatch	16
Step 3: Get Statistics from CloudWatch	17
Step 4: View Graphs with the Console	17
Using Dashboards	18
Create a Dashboard	18
Add or Remove a Graph	19
Move or Resize a Graph	20
Edit a Graph	20
Rename a Graph	21
Add or Remove a Text Widget	21
Monitor Resources in Multiple Regions	22
Link and Unlink Graphs	22
Change the Refresh Interval	23
Change the Time Range or Format	23
Using Metrics	25
View Available Metrics	25
Search for Available Metrics	28
Get Statistics for a Metric	29
Get Statistics for a Specific Resource	30
Aggregate Statistics Across Resources	33
Aggregate Statistics by Auto Scaling Group	34
Aggregate Statistics by AMI	36
Graph Metrics	37
Graph a Metric	37
Modify the Time Range for a Graph	40
Modify the Y Axis for a Graph	42
Create an Alarm from a Metric on a Graph	42
Publish Custom Metrics	43
Publish Single Data Points	43

Publish Statistic Sets	44
Publish the Value Zero	45
Metrics and Dimensions Reference	46
AWS Namespaces	47
API Gateway	48
API Gateway Metrics	49
Dimensions for Metrics	49
Auto Scaling	50
Auto Scaling Group Metrics	50
Dimensions for Auto Scaling Group Metrics	50
AWS Billing and Cost Management	51
AWS Billing and Cost Management Metrics	51
Dimensions for AWS Billing and Cost Management Metrics	51
Amazon CloudFront	51
Amazon CloudFront Metrics	51
Dimensions for CloudFront Metrics	52
Amazon CloudSearch	53
Amazon CloudSearch Metrics	53
Dimensions for Amazon CloudSearch Metrics	53
Amazon CloudWatch Events	54
CloudWatch Events Metrics	54
Dimensions for CloudWatch Events Metrics	54
Amazon CloudWatch Logs	55
CloudWatch Logs Metrics	55
Dimensions for CloudWatch Logs Metrics	56
Amazon DynamoDB	56
DynamoDB Metrics	56
Dimensions for DynamoDB Metrics	66
Amazon EC2	66
Amazon EC2 Metrics	66
Dimensions for Amazon EC2 Metrics	69
Amazon EC2 Spot Fleet	70
Amazon EC2 Spot Fleet Metrics	70
Dimensions for Amazon EC2 Spot Fleet Metrics	71
Amazon ECS	71
Amazon ECS Metrics	71
Dimensions for Amazon ECS Metrics	72
Elastic Beanstalk	73
Elastic Beanstalk Metrics	73
Dimensions for Elastic Beanstalk Metrics	74
Amazon ElastiCache	74
Dimensions for ElastiCache Metrics	75
Host-Level Metrics	75
Metrics for Memcached	76
Metrics for Redis	78
Amazon EBS	80
Amazon EBS Metrics	80
Dimensions for Amazon EBS Metrics	82
Amazon EFS	82
Amazon CloudWatch Metrics for Amazon EFS	82
Dimensions for Amazon EFS Metrics	84
Elastic Load Balancing	84
Application Load Balancer Metrics	85
Metric Dimensions for Application Load Balancers	86
Classic Load Balancer Metrics	86
Metric Dimensions for Classic Load Balancers	89
Amazon EMR	89
Amazon EMR Metrics	90

Amazon EMR Dimensions	98
Amazon ES	99
Amazon Elasticsearch Service Metrics	99
Dimensions for Amazon Elasticsearch Service Metrics	101
Elastic Transcoder	101
Elastic Transcoder Metrics	101
Dimensions for Elastic Transcoder Metrics	103
AWS IoT	103
AWS IoT Metrics	103
Dimensions for AWS IoT Metrics	104
Amazon Kinesis Analytics	105
Metrics	105
Dimensions for Metrics	105
Amazon Kinesis Firehose	105
Service-level CloudWatch Metrics	105
API-Level CloudWatch Metrics	106
Amazon Kinesis Streams	108
Basic Stream-level Metrics	108
Enhanced Shard-level Metrics	111
Dimensions for Amazon Kinesis Metrics	113
AWS KMS	114
AWS KMS Metrics	114
Dimensions for AWS KMS Metrics	114
AWS Lambda	114
CloudWatch Metrics	114
Dimensions for AWS Lambda Metrics	116
Amazon Machine Learning	116
Amazon ML Metrics	116
Dimensions for Amazon Machine Learning Metrics	116
AWS OpsWorks	117
AWS OpsWorks Metrics	117
Dimensions for AWS OpsWorks Metrics	118
Amazon Polly	118
Amazon Polly Metrics	118
Dimensions for Amazon Polly Metrics	119
Amazon Redshift	119
Amazon Redshift Metrics	119
Dimensions for Amazon Redshift Metrics	121
Amazon RDS	122
Amazon RDS Metrics	122
Dimensions for RDS Metrics	124
Amazon Route 53	124
Amazon Route 53 Metrics	124
Dimensions for Amazon Route 53 Metrics	125
Amazon SES	126
Amazon SES Event Metrics	126
Dimensions for Amazon SES Metrics	126
Amazon SNS	126
Amazon Simple Notification Service Metrics	127
Dimensions for Amazon Simple Notification Service Metrics	127
Amazon SQS	128
Amazon SQS Metrics	128
Dimensions for Amazon SQS Metrics	130
Amazon S3	130
Amazon S3 CloudWatch Metrics	130
Amazon S3 CloudWatch Dimensions	132
Amazon SWF	133
Workflow Metrics	133

Activity Metrics	134
AWS Storage Gateway	134
AWS Storage Gateway Metrics	135
Dimensions for AWS Storage Gateway Metrics	145
AWS WAF	145
AWS WAF Metrics	145
Dimensions for AWS WAF	146
Amazon WorkSpaces	146
Amazon WorkSpaces Metrics	146
Dimensions for Amazon WorkSpaces Metrics	148
Creating Alarms	149
Set Up an SNS Topic	151
Set Up an SNS Topic Using the AWS Management Console	151
Set Up an SNS Topic Using the AWS CLI	152
Create or Edit an Alarm	153
Create a CPU Usage Alarm	154
Set Up a CPU Usage Alarm Using the AWS Management Console	154
Set Up a CPU Usage Alarm Using the AWS CLI	156
Create a Load Balancer Latency Alarm	156
Set Up a Latency Alarm Using the AWS Management Console	157
Set Up a Latency Alarm Using the AWS CLI	157
Create a Storage Throughput Alarm	158
Set Up a Storage Throughput Alarm Using the AWS Management Console	158
Set Up a Storage Throughput Alarm Using the AWS CLI	158
Create Alarms to Stop, Terminate, Reboot, or Recover an Instance	159
Adding Stop Actions to Amazon CloudWatch Alarms	160
Adding Terminate Actions to Amazon CloudWatch Alarms	161
Adding Reboot Actions to Amazon CloudWatch Alarms	162
Adding Recover Actions to Amazon CloudWatch Alarms	163
Viewing the History of Triggered Alarms and Actions	165
Create a Billing Alarm	165
Enable Billing Alerts	165
Create a Billing Alarm	166
Check the Alarm Status	167
Delete a Billing Alarm	168
Authentication and Access Control	169
Authentication	169
Access Control	170
Overview of Managing Access	171
Resources and Operations	171
Understanding Resource Ownership	173
Managing Access to Resources	173
Specifying Policy Elements: Actions, Effects, and Principals	174
Specifying Conditions in a Policy	174
Using Identity-Based Policies (IAM Policies)	175
Permissions Required to Use the CloudWatch Console	175
AWS Managed (Predefined) Policies for CloudWatch	178
Customer Managed Policy Examples	178
Amazon CloudWatch Permissions Reference	180
Logging API Calls	186
CloudWatch Information in CloudTrail	186
Understanding Log File Entries	188
Document History	191

What is Amazon CloudWatch?

Amazon CloudWatch monitors your Amazon Web Services (AWS) resources and the applications you run on AWS in real time. You can use CloudWatch to collect and track metrics, which are variables you can measure for your resources and applications. CloudWatch alarms send notifications or automatically make changes to the resources you are monitoring based on rules that you define. For example, you can monitor the CPU usage and disk reads and writes of your Amazon EC2 instances and then use this data to determine whether you should launch additional instances to handle increased load. You can also use this data to stop under-used instances to save money. In addition to monitoring the built-in metrics that come with AWS, you can monitor your own custom metrics. With CloudWatch, you gain system-wide visibility into resource utilization, application performance, and operational health.

Accessing CloudWatch

You can access CloudWatch using any of the following methods:

- **Amazon CloudWatch console** — <https://console.aws.amazon.com/cloudwatch/>
- **AWS CLI** — For more information, see [Getting Set Up with the AWS Command Line Interface](#) in the *AWS Command Line Interface User Guide*.
- **CloudWatch API** — For more information, see the [Amazon CloudWatch API Reference](#).
- **AWS SDKs** — For more information, see [Tools for Amazon Web Services](#).

Related AWS Services

The following services are used in conjunction with Amazon CloudWatch:

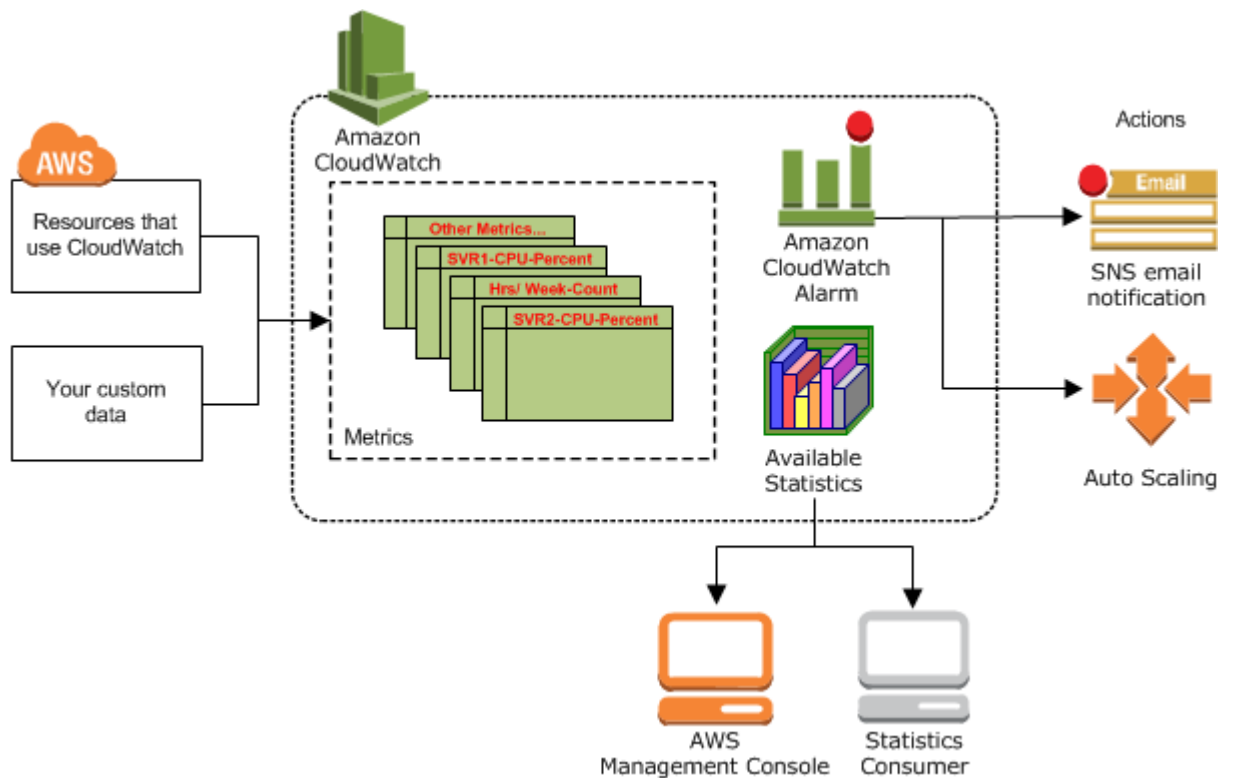
- **Amazon Simple Notification Service (Amazon SNS)** coordinates and manages the delivery or sending of messages to subscribing endpoints or clients. You use Amazon SNS with CloudWatch to send messages when an alarm threshold has been reached. For more information, see [Set Up Amazon SNS Notifications \(p. 151\)](#).
- **Auto Scaling** enables you to automatically launch or terminate Amazon EC2 instances based on user-defined policies, health status checks, and schedules. You can use a CloudWatch alarm with Auto Scaling to scale your EC2 instances based on demand. For more information, see [Dynamic Scaling](#) in the *Auto Scaling User Guide*.
- **AWS CloudTrail** enables you to monitor the calls made to the Amazon CloudWatch API for your account, including calls made by the AWS Management Console, AWS CLI, and other services.

When CloudTrail logging is turned on, CloudWatch writes log files to the Amazon S3 bucket that you specified when you configured CloudTrail. For more information, see [Logging Amazon CloudWatch API Calls in AWS CloudTrail](#) (p. 186).

- **AWS Identity and Access Management (IAM)** is a web service that helps you securely control access to AWS resources for your users. Use IAM to control who can use your AWS resources (authentication) and what resources they can use in which ways (authorization). For more information, see [Authentication and Access Control for Amazon CloudWatch](#) (p. 169).

How Amazon CloudWatch Works

Amazon CloudWatch is basically a metrics repository. An AWS product—such as Amazon EC2—puts metrics into the repository, and you retrieve statistics based on those metrics. If you put your own custom metrics into the repository, you can retrieve statistics on these metrics as well.



You can use metrics to calculate statistics and then present the data graphically in the CloudWatch console. For more information about the other AWS resources that generate and send metrics to CloudWatch, see [Amazon CloudWatch Metrics and Dimensions Reference](#) (p. 46).

You can configure alarm actions to stop, start, or terminate an Amazon EC2 instance when certain criteria are met. In addition, you can create alarms that initiate Auto Scaling and Amazon Simple Notification Service (Amazon SNS) actions on your behalf. For more information about creating CloudWatch alarms, see [Alarms](#) (p. 7).

Amazon cloud computing resources are housed in highly available data center facilities. To provide additional scalability and reliability, each data center facility is located in a specific geographical area, known as a *region*. Each region is designed to be completely isolated from the other regions, to achieve the greatest possible failure isolation and stability. Amazon CloudWatch does not aggregate data across regions. Therefore, metrics are completely separate between regions. For more information, see [Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

Amazon CloudWatch Concepts

The following terminology and concepts are central to your understanding and use of Amazon CloudWatch:

- [Namespaces \(p. 3\)](#)
- [Metrics \(p. 3\)](#)
- [Dimensions \(p. 4\)](#)
- [Statistics \(p. 5\)](#)
- [Percentiles \(p. 7\)](#)
- [Alarms \(p. 7\)](#)

Namespaces

A *namespace* is a container for CloudWatch metrics. Metrics in different namespaces are isolated from each other, so that metrics from different applications are not mistakenly aggregated into the same statistics.

There is no default namespace. You must specify a namespace for each data point you publish to CloudWatch. You can specify a namespace name when you create a metric. These names must contain valid XML characters, and be fewer than 256 characters in length. Possible characters are: alphanumeric characters (0-9A-Za-z), period (.), hyphen (-), underscore (_), forward slash (/), hash (#), and colon (:).

The AWS namespaces use the following naming convention: `AWS/service`. For example, Amazon EC2 uses the `AWS/EC2` namespace. For the list of AWS namespaces, see [AWS Namespaces \(p. 47\)](#).

Metrics

Metrics are the fundamental concept in CloudWatch. A metric represents a time-ordered set of data points that are published to CloudWatch. Think of a metric as a variable to monitor, and the data points represent the values of that variable over time. For example, the CPU usage of a particular EC2 instance is one metric provided by Amazon EC2. The data points themselves can come from any application or business activity from which you collect data.

AWS services send metrics to CloudWatch, and you can send your own custom metrics to CloudWatch. You can add the data points in any order, and at any rate you choose. You can retrieve statistics about those data points as an ordered set of time-series data.

Metrics exist only in the region in which they are created. Metrics cannot be deleted, but they automatically expire after 15 months if no new data is published to them. Data points older than 15 months expire on a rolling basis; as new data points come in, data older than 15 months is dropped.

Metrics are uniquely defined by a name, a namespace, and one or more dimensions. Each data point has a time stamp, and (optionally) a unit of measure. When you request statistics, the returned data stream is identified by namespace, metric name, dimension, and (optionally) the unit.

For more information, see [View Available Metrics \(p. 25\)](#) and [Publish Custom Metrics \(p. 43\)](#).

Time Stamps

Each metric data point must be marked with a time stamp. The time stamp can be up to two weeks in the past and up to two hours into the future. If you do not provide a time stamp, CloudWatch creates a time stamp for you based on the time the data point was received.

Time stamps are `dateTime` objects, with the complete date plus hours, minutes, and seconds (for example, 2016-10-31T23:59:59Z). For more information, see [dateTime](#). Although it is not required, we recommend that you use Coordinated Universal Time (UTC). When you retrieve statistics from CloudWatch, all times are in UTC.

Note that CloudWatch alarms check metrics based on the current time in UTC. Custom metrics sent to CloudWatch with time stamps other than the current UTC time can cause alarms to display the **Insufficient Data** state or result in delayed alarms.

Metrics Retention

CloudWatch retains your metric data as follows:

- Data points with a period of 60 seconds (1 minute) are available for 15 days
- Data points with a period of 300 seconds (5 minute) are available for 63 days
- Data points with a period of 3600 seconds (1 hour) are available for 455 days (15 months)

Note that CloudWatch started retaining 5-minute and 1-hour metric data as of 9 July 2016.

Dimensions

A *dimension* is a name/value pair that uniquely identifies a metric. You can assign up to ten dimensions to a metric.

Every metric has specific characteristics that describe it, and you can think of dimensions as categories for those characteristics. Dimensions help you design a structure for your statistics plan. Because dimensions are part of the unique identifier for a metric, whenever you add a unique name/value pair to one of your metrics, you are creating a new metric.

AWS services that send data to CloudWatch attach dimensions to each metric. You can use dimensions to filter the results that CloudWatch returns. For example, you can get statistics for a specific EC2 instance by specifying the `InstanceId` dimension when you search for metrics.

For metrics produced by certain AWS services, such as Amazon EC2, CloudWatch can aggregate data across dimensions. For example, search for metrics in the `AWS/EC2` namespace but do not specify any dimensions, CloudWatch aggregates all data for the specified metric to create the statistic that you requested. Note that CloudWatch does not aggregate across dimensions for your custom metrics.

Dimension Combinations

CloudWatch treats each unique combination of dimensions as a separate metric, even if the metrics have the same metric name. You can't retrieve statistics using combinations of dimensions that you did not specifically publish. When you retrieve statistics, specify the same values for the namespace, metric name, and dimension parameters that were used when the metrics were created. You can also specify the start and end times for CloudWatch to use for aggregation.

For example, suppose that you publish four distinct metrics named `ServerStats` in the `DataCenterMetric` namespace with the following properties:

```
Dimensions: Server=Prod, Domain=Frankfurt, Unit: Count, Timestamp:
 2016-10-31T12:30:00Z, Value: 105
Dimensions: Server=Beta, Domain=Frankfurt, Unit: Count, Timestamp:
 2016-10-31T12:31:00Z, Value: 115
Dimensions: Server=Prod, Domain=Rio, Unit: Count, Timestamp:
 2016-10-31T12:32:00Z, Value: 95
Dimensions: Server=Beta, Domain=Rio, Unit: Count, Timestamp:
 2016-10-31T12:33:00Z, Value: 97
```

If you publish only those four metrics, you can retrieve statistics for these combinations of dimensions:

- Server=Prod,Domain=Frankfurt
- Server=Prod,Domain=Rio
- Server=Beta,Domain=Frankfurt
- Server=Beta,Domain=Rio

You can't retrieve statistics for the following dimensions or if you specify no dimensions:

- Server=Prod
- Server=Beta
- Domain=Frankfurt
- Domain=Rio

Statistics

Statistics are metric data aggregations over specified periods of time. CloudWatch provides statistics based on the metric data points provided by your custom data or provided by other services in AWS to CloudWatch. Aggregations are made using the namespace, metric name, dimensions, and the data point unit of measure, within the time period you specify. The following table describes the available statistics.

Statistic	Description
Minimum	The lowest value observed during the specified period. You can use this value to determine low volumes of activity for your application.
Maximum	The highest value observed during the specified period. You can use this value to determine high volumes of activity for your application.
Sum	All values submitted for the matching metric added together. This statistic can be useful for determining the total volume of a metric.
Average	The value of <code>Sum / SampleCount</code> during the specified period. By comparing this statistic with the <code>Minimum</code> and <code>Maximum</code> , you can determine the full scope of a metric and how close the average use is to the <code>Minimum</code> and <code>Maximum</code> . This comparison helps you to know when to increase or decrease your resources as needed.
SampleCount	The count (number) of data points used for the statistical calculation.
pNN.NN	The value of the specified percentile. You can specify any percentile, using up to two decimal places (for example, p95.45). For more information, see Percentiles (p. 7) .

You can add pre-calculated statistics. Instead of data point values, you specify values for `SampleCount`, `Minimum`, `Maximum`, and `Sum` (CloudWatch calculates the average for you). The values you add in this way are aggregated with any other values associated with the matching metric.

Units

Each statistic has a unit of measure. Example units include `Bytes`, `Seconds`, `Count`, and `Percent`. For the complete list of the units that CloudWatch supports, see the [MetricDatum](#) data type in the *Amazon CloudWatch API Reference*.

You can specify a unit when you create a custom metric. If you do not specify a unit, CloudWatch uses `None` as the unit. Units help provide conceptual meaning to your data. Though CloudWatch attaches no significance to a unit internally, other applications can derive semantic information based on the unit.

Metric data points that specify a unit of measure are aggregated separately. When you get statistics without specifying a unit, CloudWatch aggregates all data points of the same unit together. If you have two otherwise identical metrics with different units, two separate data streams are returned, one for each unit.

Periods

A *period* is the length of time associated with a specific Amazon CloudWatch statistic. Each statistic represents an aggregation of the metrics data collected for a specified period of time. Although periods are expressed in seconds, the minimum granularity for a period is one minute. Accordingly, you specify period values as multiples of 60. For example, to specify a period of six minutes, you would use the value 360. You can adjust how the data is aggregated by varying the length of the period. A period can be as short as one minute (60 seconds) or as long as one day (86,400 seconds). The default value is 60 seconds.

When you retrieve statistics, you can specify a period, start time, and end time. These parameters determine the overall length of time associated with the statistics. The default values for the start time and end time get you the last hour's worth of statistics. The values that you specify for the start time and end time determine how many periods CloudWatch will return. For example, retrieving statistics using the default values for the period, start time, and end time returns an aggregated set of statistics for each minute of the previous hour. If you prefer statistics aggregated in ten-minute blocks, specify a period of 600. For statistics aggregated over the entire hour, specify a period of 3600.

Periods are also important for CloudWatch alarms. When you create an alarm to monitor a specific metric, you are asking CloudWatch to compare that metric to the threshold value that you specified. You have extensive control over how CloudWatch makes that comparison. Not only can you specify the period over which the comparison is made, but you can also specify how many evaluation periods are used to arrive at a conclusion. For example, if you specify three evaluation periods, CloudWatch compares a window of three data points. CloudWatch only notifies you if the oldest data point is breaching and the others are breaching or missing. For metrics that are continuously emitted, CloudWatch won't notify you until three failures are found.

Aggregation

Amazon CloudWatch aggregates statistics according to the period length that you specify when retrieving statistics. You can publish as many data points as you want with the same or similar time stamps. CloudWatch aggregates them by period length. Aggregated statistics are only available when using detailed monitoring. In addition, Amazon CloudWatch does not aggregate data across regions.

You can publish data points for a metric that share not only the same time stamp, but also the same namespace and dimensions. CloudWatch will return aggregated statistics for those data points. You can also publish multiple data points for the same or different metrics, with any time stamp.

For large data sets, you can insert a pre-aggregated data set called a *statistic set*. With statistic sets, you give CloudWatch the Min, Max, Sum, and SampleCount for a number of data points. This is commonly used when you need to collect data many times in a minute. For example, suppose you have a metric for the request latency of a web page. It doesn't make sense to publish data with every web page hit. We suggest that you collect the latency of all hits to that web page, aggregate them once a minute, and send that statistic set to CloudWatch.

Amazon CloudWatch doesn't differentiate the source of a metric. If you publish a metric with the same namespace and dimensions from different sources, CloudWatch treats this as a single metric. This can be useful for service metrics in a distributed, scaled system. For example, all the hosts in a web server application could publish identical metrics representing the latency of requests they are processing.

CloudWatch treats these as a single metric, allowing you to get the statistics for minimum, maximum, average, and sum of all requests across your application.

Percentiles

A *percentile* indicates the relative standing of a value in a data set. For example, the 95th percentile means that 95 percent of the data is below this value and 5 percent of the data is above this value. Percentiles help you get a better understanding of the distribution of your metric data.

Percentiles are often used to isolate anomalies. In a typical distribution, 95 percent of the data is within two standard deviations from the mean and 99.7 percent of the data is within three standard deviations from the mean. Any data that falls outside three standard deviations is often considered to be an anomaly because it differs so greatly from the average value. For example, suppose that you are monitoring the CPU utilization of your EC2 instances to ensure that your customers have a good experience. If you monitor the average, this can hide anomalies. If you monitor the maximum, a single anomaly can skew the results. Using percentiles, you can monitor the 95th percentile of CPU utilization to check for instances with an unusually heavy load.

You can monitor your system and applications using percentiles as you would use the other CloudWatch statistics (Average, Minimum, Maximum, and Sum). For example, when you create an alarm, you can use percentiles as the statistical function. You can specify the percentile with up to two decimal places (for example, p95.45).

Alarms

You can use an *alarm* to automatically initiate actions on your behalf. An alarm watches a single metric over a specified time period, and performs one or more specified actions, based on the value of the metric relative to a threshold over time. The action is a notification sent to an Amazon SNS topic or an Auto Scaling policy.

Alarms invoke actions for sustained state changes only. CloudWatch alarms will not invoke actions simply because they are in a particular state. The state must have changed and been maintained for a specified number of periods.

When creating an alarm, select a period that is greater than or equal to the frequency of the metric to be monitored. For example, basic monitoring for Amazon EC2 provides metrics for your instances every 5 minutes. When setting an alarm on a basic monitoring metric, select a period of at least 300 seconds (5 minutes). Detailed monitoring for Amazon EC2 provides metrics for your instances every 1 minute. When setting an alarm on a detailed monitoring metric, select a period of at least 60 seconds (1 minute).

For more information, see [Creating Amazon CloudWatch Alarms \(p. 149\)](#) and [Create an Alarm from a Metric on a Graph \(p. 42\)](#).

CloudWatch Limits

CloudWatch has the following limits:

Resource	Default Limit
Actions	5/alarm. This limit cannot be changed.
Alarms	10/month/customer for free. 5000/account.
API requests	1,000,000/month/customer for free.
Custom metrics	No limit.

Resource	Default Limit
DescribeAlarms	3 transactions per second (TPS). The maximum number of operation requests you can make per second without being throttled. You can request a limit increase .
Dimensions	10/metric. This limit cannot be changed.
GetMetricStatistics	400 transactions per second (TPS). The maximum number of operation requests you can make per second without being throttled. You can request a limit increase .
ListMetrics	25 transactions per second (TPS). The maximum number of operation requests you can make per second without being throttled. You can request a limit increase .
Metric data	15 months. This limit cannot be changed.
MetricDatum items	20/ PutMetricData request. A MetricDatum object can contain a single value or a StatisticSet object representing many values. This limit cannot be changed.
Metrics	10/month/customer for free.
Period	One day (86,400 seconds). This limit cannot be changed.
PutMetricAlarm request	3 transactions per second (TPS). The maximum number of operation requests you can make per second without being throttled. You can request a limit increase .
PutMetricData request	8 KB for HTTP GET requests and 40 KB for HTTP POST requests. PutMetricData can handle 150 transactions per second (TPS), which is the maximum number of operation requests you can make per second without being throttled. You can request a limit increase .
Amazon SNS email notifications	1,000/month/customer for free.

Amazon CloudWatch Resources

The following table lists related resources that you'll find useful as you work with Amazon CloudWatch.

Resource	Description
Amazon CloudWatch FAQs	The FAQ covers the top questions developers have asked about this product.
Release notes	The release notes give a high-level overview of the current release. They specifically note any new features, corrections, and known issues.

Resource	Description
AWS Developer Resource Center	A central starting point to find documentation, code samples, release notes, and other information to help you build innovative applications with AWS.
AWS Management Console	The console allows you to perform most of the functions of Amazon CloudWatch and various other AWS products without programming.
Amazon CloudWatch Discussion Forums	Community-based forum for developers to discuss technical questions related to Amazon CloudWatch.
AWS Support	The hub for creating and managing your AWS Support cases. Also includes links to other helpful resources, such as forums, technical FAQs, service health status, and AWS Trusted Advisor.
Amazon CloudWatch product information	The primary web page for information about Amazon CloudWatch.
Contact Us	A central contact point for inquiries concerning AWS billing, account, events, abuse, etc.

Getting Set Up

To use Amazon CloudWatch you need an AWS account. Your AWS account allows you to use services (for example, Amazon EC2) to generate metrics that you can view in the CloudWatch console, a point-and-click web-based interface. In addition, you can install and configure the AWS command line interface (CLI).

Sign Up for Amazon Web Services (AWS)

When you create an AWS account, we automatically sign up your account for all AWS services. You pay only for the services that you use.

If you have an AWS account already, skip to the next step. If you don't have an AWS account, use the following procedure to create one.

To sign up for an AWS account

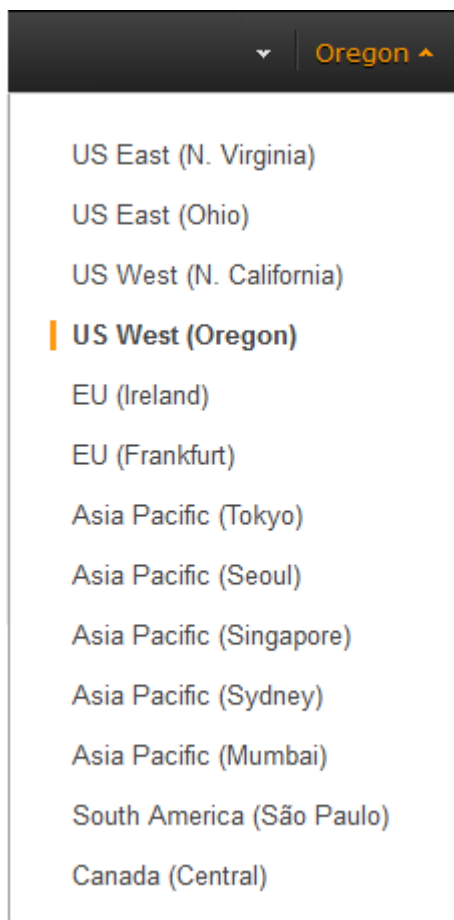
1. Open <http://aws.amazon.com/>, and then choose **Create an AWS Account**.
2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call and entering a PIN using the phone keypad.

Sign in to the Amazon CloudWatch Console

To sign in to the Amazon CloudWatch console

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. If necessary, use the navigation bar to change the region to the region where you have your AWS resources.



3. Even if this is the first time you are using the CloudWatch console, **Your Metrics** could already report metrics, because you have used a AWS product that automatically pushes metrics to Amazon CloudWatch for free. Other AWS products require that you enable metrics.

If you do not have any alarms, the **Your Alarms** section will have a **Create Alarm** button.

Set Up the AWS CLI

You can use the AWS CLI or the Amazon CloudWatch CLI to perform CloudWatch commands. Note that the AWS CLI replaces the CloudWatch CLI; we include new CloudWatch features only in the AWS CLI.

For information about how to install and configure the AWS CLI, see [Getting Set Up with the AWS Command Line Interface](#) in the *AWS Command Line Interface User Guide*.

For information about how to install and configure the Amazon CloudWatch CLI, see [Set Up the Command Line Interface](#) in the *Amazon CloudWatch CLI Reference*.

Getting Started with Amazon CloudWatch

The following scenarios show you how to use Amazon CloudWatch. In the first scenario, you use the CloudWatch console to create a billing alarm that tracks your AWS usage and lets you know when you have exceeded a certain spending threshold. In the second, more advanced scenario, you use the AWS command line interface (CLI) to publish a single metric for a hypothetical application named *GetStarted*.

Scenarios

- [Monitor Your Estimated Charges \(p. 12\)](#)
- [Publish Metrics \(p. 15\)](#)

Scenario: Monitor Your Estimated Charges Using CloudWatch

In this scenario, you create an Amazon CloudWatch alarm to monitor your estimated charges. When you enable the monitoring of estimated charges for your AWS account, the estimated charges are calculated and sent several times daily to CloudWatch as metric data.

Billing metric data is stored in the US East (N. Virginia) Region and reflects worldwide charges. This data includes the estimated charges for every service in AWS that you use, as well as the estimated overall total of your AWS charges.

You can choose to receive alerts by email when charges have exceeded a certain threshold. These alerts are triggered by CloudWatch and messages are sent using Amazon Simple Notification Service (Amazon SNS).

Tasks

- [Step 1: Enable Billing Alerts \(p. 13\)](#)
- [Step 2: Create a Billing Alarm \(p. 13\)](#)
- [Step 3: Check the Alarm Status \(p. 14\)](#)

- [Step 4: Edit a Billing Alarm \(p. 15\)](#)
- [Step 5: Delete a Billing Alarm \(p. 15\)](#)

Step 1: Enable Billing Alerts

Before you can create an alarm for your estimated charges, you must enable billing alerts, so that you can monitor your estimated AWS charges and create an alarm using billing metric data. After you enable billing alerts, you cannot disable data collection, but you can delete any billing alarms you created.

After you enable billing alerts for the first time, it takes about 15 minutes before you can view billing data and set billing alarms.

Requirements

- You must be signed in using root account credentials; IAM users cannot enable billing alerts for your AWS account.
- For consolidated billing accounts, billing data for each linked account can be found by logging in as the paying account. You can view billing data for total estimated charges and estimated charges by service for each linked account as well as for the consolidated account.

To enable monitoring of your estimated charges

1. Open the Billing and Cost Management console at <https://console.aws.amazon.com/billing/home?#>.
2. In the navigation pane, choose **Preferences**.
3. Select **Receive Billing Alerts**.

The screenshot shows the 'Preferences' page in the AWS Billing and Cost Management console. On the left is a navigation menu with options: Dashboard, Bills, Cost Explorer, Budgets, Reports, Cost Allocation Tags, Payment Methods, Payment History, Consolidated Billing, Preferences (highlighted), Credits, Tax Settings, and DevPay. The main content area is titled 'Preferences' and contains three settings:

- Receive PDF Invoice By Email**: Turn on this feature to receive a PDF version of your invoice by email. Invoices are generally available within the first three days of the month.
- Receive Billing Alerts**: Turn on this feature to monitor your AWS usage charges and recurring fees automatically, making it easier to track and manage your spending on AWS. You can set up billing alerts to receive email notifications when your charges reach a specified threshold. Once enabled, this preference cannot be disabled. [Manage Billing Alerts](#)
- Receive Billing Reports**: Turn on this feature to receive ongoing reports of your AWS charges once or more daily. AWS delivers these reports to the Amazon S3 bucket that you specify where indicated below. For consolidated billing customers, AWS generates reports only for paying accounts. Linked accounts cannot sign up for billing reports.

Below the 'Receive Billing Reports' section, there is a 'Save to S3 Bucket:' label, a text input field containing 'bucket name', and a 'Verify' button. At the bottom of the main content area is a blue 'Save preferences' button.

4. Choose **Save preferences**.

Step 2: Create a Billing Alarm

After you've enabled billing alerts, you can create a billing alarm. In this scenario, you create an alarm that sends an email message when your estimated charges for AWS exceed a specified threshold.

Note

This procedure uses the simple options. To use the advanced options, see [Create a Billing Alarm](#) (p. 166) in *Create a Billing Alarm to Monitor Your Estimated AWS Charges*.

To create a billing alarm

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. If necessary, change the region to US East (N. Virginia). Billing metric data is stored in this region and reflects worldwide charges.
3. In the navigation pane, choose **Alarms**, **Billing**.
4. For **Whenever my total AWS charges for the month exceed**, specify the monetary amount (for example, 200) that must be exceeded to trigger the alarm and send an email notification.

Tip

Under **Alarm Preview**, there is an estimate of your charges that you can use to set an appropriate amount.

The screenshot shows the 'Create Alarm' wizard in the AWS CloudWatch console. The main heading is 'Billing Alarm'. Below it, there is a brief explanation and a numbered list of steps: 1. Enter a spending threshold, 2. Provide an email address, and 3. Check your inbox for a confirmation email and click the link provided. A form field is set to 'When my total AWS charges for the month exceed: \$ 200 USD'. Below this, there is a dropdown menu for 'send a notification to:' with a 'New list' link. A reminder note states that users will receive an email from AWS with the subject 'AWS Notification - Subscription Confirmation' and should click the link to confirm delivery. At the bottom of the form, it says 'showing simple options | show advanced'. On the right side, the 'Alarm Preview' section shows a line graph titled 'EstimatedCharges >= 200'. The graph has a y-axis from 0 to 250 and an x-axis with dates 10/25 00:00, 10/27 00:00, and 10/29 00:00. A red horizontal line is drawn at the 200 mark. A blue area under the curve shows the estimated charges, which are below the red line. Below the graph, there are links for 'More resources': 'AWS Billing console', 'Getting started with billing alarms', 'More help with billing alarms', and 'AWS Billing FAQs'. At the bottom of the console window, there are buttons for 'Cancel', 'Previous', 'Next', and 'Create Alarm'.

5. For **send a notification to**, choose an existing notification list or create a new one.
To create a list, choose **New list** and type a comma-separated list of email addresses to be notified when the alarm changes to the ALARM state. Each email address will be sent a subscription confirmation email. The recipient must confirm the subscription before notifications can be sent to the email address.
6. Choose **Create Alarm**.

Step 3: Check the Alarm Status

Now, check the status of the billing alarm that you just created.

To check the alarm status

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.

2. If necessary, change the region to US East (N. Virginia). Billing metric data is stored in this region and reflects worldwide charges.
3. In the navigation pane, choose **Alarms, Billing**.
4. Select the check box next to the alarm. Note that until the subscription is confirmed, it is shown as "Pending confirmation". After the subscription is confirmed, refresh the console to show the updated status.

Step 4: Edit a Billing Alarm

Let's say that you want to increase the amount money you spend with AWS each month from \$200 to \$400. You can edit your existing billing alarm and increase the monetary amount that must be exceeded before the alarm is triggered.

To edit a billing alarm

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. If necessary, change the region to US East (N. Virginia). Billing metric data is stored in this region and reflects worldwide charges.
3. In the navigation pane, choose **Alarms, Billing**.
4. Select the check box next to the alarm and then choose **Modify**.
5. For **Whenever my total AWS charges for the month exceed**, specify the new amount that must be exceeded to trigger the alarm and send an email notification.
6. Choose **Save Changes**.

Step 5: Delete a Billing Alarm

You can delete your billing alarm if you no longer need it.

To delete a billing alarm

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. If necessary, change the region to US East (N. Virginia). Billing metric data is stored in this region and reflects worldwide charges.
3. In the navigation pane, choose **Alarms, Billing**.
4. Select the check box next to the alarm and then choose **Delete**.
5. When prompted for confirmation, choose **Yes, Delete**.

Scenario: Publish Metrics to CloudWatch

In this scenario, you'll use the AWS Command Line Interface (AWS CLI) to publish a single metric for a hypothetical application named *GetStarted*. If you haven't already installed and configured the AWS CLI, see [Getting Set Up with the AWS Command Line Interface](#) in the *AWS Command Line Interface User Guide*.

Tasks

- [Step 1: Define the Data Configuration](#) (p. 16)
- [Step 2: Add Metrics to CloudWatch](#) (p. 16)
- [Step 3: Get Statistics from CloudWatch](#) (p. 17)
- [Step 4: View Graphs with the Console](#) (p. 17)

Step 1: Define the Data Configuration

In this scenario, you'll publish data points that track the request latency for the application. Choose names for your metric and namespace that make sense to you. For this example, name the metric *RequestLatency* and place all of the data points into the *GetStarted* namespace.

You'll publish several data points that collectively represent three hours of latency data. The raw data comprises fifteen request latency readings distributed over three hours. Each reading is in milliseconds:

- Hour one: 87, 51, 125, 235
- Hour two: 121, 113, 189, 65, 89
- Hour three: 100, 47, 133, 98, 100, 328

You can publish data to CloudWatch as single data points or as an aggregated set of data points called a *statistic set*. You can aggregate metrics to a granularity as low as one minute. You can publish the aggregated data points to CloudWatch as a set of statistics with four predefined keys: *Sum*, *Minimum*, *Maximum*, and *SampleCount*.

You'll publish the data points from hour one as single data points. For the data from hours two and three, you'll aggregate the data points and publish a statistic set for each hour. The key values are shown in the following table.

Hour	Raw Data	Sum	Minimum	Maximum	SampleCount
1	87				
1	51				
1	125				
1	235				
2	121, 113, 189, 65, 89	577	65	189	5
3	100, 47, 133, 98, 100, 328	806	47	328	6

Step 2: Add Metrics to CloudWatch

After you have defined your data configuration, you are ready to add data.

To publish data points to CloudWatch

1. At a command prompt, run the following `put-metric-data` commands to add data for the first hour. Replace the example time stamp with a time stamp that is two hours in the past, in Universal Coordinated Time (UTC).

```
aws cloudwatch put-metric-data --metric-name RequestLatency --namespace
  GetStarted \
--timestamp 2016-10-14T20:30:00Z --value 87 --unit Milliseconds
aws cloudwatch put-metric-data --metric-name RequestLatency --namespace
  GetStarted \
--timestamp 2016-10-14T20:30:00Z --value 51 --unit Milliseconds
aws cloudwatch put-metric-data --metric-name RequestLatency --namespace
  GetStarted \
--timestamp 2016-10-14T20:30:00Z --value 125 --unit Milliseconds
```

```
aws cloudwatch put-metric-data --metric-name RequestLatency --namespace
  GetStarted \
--timestamp 2016-10-14T20:30:00Z --value 235 --unit Milliseconds
```

2. Add data for the second hour, using a time stamp that is one hour later than the first hour.

```
aws cloudwatch put-metric-data --metric-name RequestLatency --namespace
  GetStarted \
--timestamp 2016-10-14T21:30:00Z --statistic-values
  Sum=577,Minimum=65,Maximum=189,SampleCount=5 --unit Milliseconds
```

3. Add data for the third hour, omitting the time stamp to default to the current time.

```
aws cloudwatch put-metric-data --metric-name RequestLatency --namespace
  GetStarted \
--statistic-values Sum=806,Minimum=47,Maximum=328,SampleCount=6 --unit
  Milliseconds
```

Step 3: Get Statistics from CloudWatch

Now that you have published metrics to CloudWatch, you can retrieve statistics based on those metrics using the `get-metric-statistics` command as follows. Be sure to specify `--start-time` and `--end-time` far enough in the past to cover the earliest time stamp that you published.

```
aws cloudwatch get-metric-statistics --namespace GetStarted --metric-name
  RequestLatency --statistics Average \
--start-time 2016-10-14T00:00:00Z --end-time 2016-10-15T00:00:00Z --period 60
```

The following is example output:

```
{
  "Datapoints": [],
  "Label": "Request:Latency"
}
```

Step 4: View Graphs with the Console

After you have published metrics to CloudWatch, you can use the CloudWatch console to view statistical graphs.

To view graphs of your statistics on the console

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the **Navigation** pane, choose **Metrics**.
3. In the **All metrics** tab, in the search box, type **RequestLatency** and press Enter.
4. Select the check box for the **RequestLatency** metric. A graph of the metric data is displayed in the upper pane.

For more information, see [Graph Metrics \(p. 37\)](#).

Using Amazon CloudWatch Dashboards

Amazon CloudWatch dashboards are customizable home pages in the CloudWatch console that you can use to monitor your resources in a single view, even those resources that are spread across different regions. You can use CloudWatch dashboards to create customized views of the metrics for your AWS resources.

With dashboards, you can create the following:

- A single view for selected metrics to help you assess the health of your resources and applications across one or more regions.
- An operational playbook that provides guidance for team members during operational events about how to respond to specific incidents.
- A common view of critical resource and application measurements that can be shared by team members for faster communication flow during operational events.

Contents

- [Create a Dashboard \(p. 18\)](#)
- [Add or Remove a Graph \(p. 19\)](#)
- [Move or Resize a Graph \(p. 20\)](#)
- [Edit a Graph \(p. 20\)](#)
- [Rename a Graph \(p. 21\)](#)
- [Add or Remove a Text Widget \(p. 21\)](#)
- [Monitor Resources in Multiple Regions \(p. 22\)](#)
- [Link and Unlink Graphs \(p. 22\)](#)
- [Change the Refresh Interval \(p. 23\)](#)
- [Change the Time Range or Format \(p. 23\)](#)

Create a CloudWatch Dashboard

To get started with CloudWatch dashboards, you must first create a dashboard. Note that you can create multiple dashboards to track metrics for your AWS resources.

To create a dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose **Create dashboard**.
4. In the **Create new dashboard** dialog box, type a name for the dashboard and then choose **Create dashboard**.
5. Do one of the following in the **Add widget to dashboard** dialog box:
 - To add a graph to your dashboard, choose **Metric graph** and then choose **Configure**. Then, in the **Add metric graph** dialog box, select the metrics to graph, and then choose **Create widget**.
 - To add a text block to your dashboard, choose **Text widget** and then choose **Configure**. Then, in the **New text widget** dialog box, for **Markdown**, add and format your text using [Markdown](#), and then choose **Create widget**.
6. Choose **Save dashboard**.

Add or Remove a Graph from a CloudWatch Dashboard

You can add graphs containing one or more metrics to your dashboard for the resources you monitor. You can remove the graphs when they're no longer needed.

To add a graph to a dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose your dashboard and then choose **Add widget**.
4. In the **Add widget to dashboard** dialog box, choose a widget type (**Line**, **Stacked area**, **Number**, or **Text**) and then choose **Configure**.
5. Choose the metrics to graph and then choose **Create widget**.
6. (Optional) To temporarily make the graph larger, select the graph.
7. (Optional) To view more information about the metric being graphed, hover over the legend.
8. (Optional) To change the widget type, hover over the title of the graph and choose **Widget actions**, **Widget type**.
9. Choose **Save dashboard**.

To add a graph from an alarm to a dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Alarms**.
3. Select an alarm. On the **Details** tab, select the graph
4. Choose **Actions**, **Add to dashboard**.
5. In the **Add to dashboard** dialog box, for **Add to**, choose a dashboard, and then choose **Add to dashboard**.
6. Choose **Save dashboard**.

To remove a graph from a dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.

2. In the navigation pane, choose **Dashboards**.
3. Choose your dashboard.
4. Hover over the title of the graph, choose **Widget actions**, **Delete**.
5. Choose **Save dashboard**. Note that if you attempt to navigate away from the dashboard before you save your changes, you are prompted to either save or discard your changes.

Move or Resize a Graph on a CloudWatch Dashboard

You can arrange and resize graphs on your CloudWatch dashboard.

To move or resize a graph on a dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose a dashboard.
4. Hover over the title of the graph until the selection icon appears, and then select and drag the graph to a new location on the dashboard.
5. Choose **Save dashboard**.

To resize a graph

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose a dashboard.
4. Hover over the graph, then drag the lower right corner of the graph to increase or decrease the size.
5. Choose **Save dashboard**.

To enlarge a graph temporarily

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose a dashboard.
4. Select the graph. Alternatively, hover over the title of the graph and choose **Widget actions**, **Enlarge**.

Edit a Graph on a CloudWatch Dashboard

You can edit a graph to change the title, statistic, or period, or to add or remove metrics. If you have multiple metrics displayed on a graph, you can reduce clutter by temporarily hiding the metrics that don't interest you.

To edit a graph on a dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.

3. In the **Dashboards** list, choose a dashboard.
4. Hover over the title of the graph, choose **Widget actions, Edit**.
5. In the lower half of the **Edit Graph** screen, you can change the title, statistic, or period:
 - a. To change the graph's title, select the title, type a new title, and then choose **Save**.
 - b. To change the statistic, which is next to the graph's title, choose **Statistic**, and then choose another value.
 - c. To change the time period, which is next to **Statistic**, choose **Period**, and then choose another value.
6. When you're finished with your changes, choose **Update widget**.

To temporarily hide metrics on a graph on a dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. In the **Dashboards** list, choose a dashboard.
4. In the graph's footer, hover over the colored square in the legend, and then when it changes to an X, click it.
5. To restore the metric, click the grayed out square and metric name.

Rename a Graph on a CloudWatch Dashboard

You can change the default name that CloudWatch assigns to a graph on your dashboard.

To rename a graph on a dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose a dashboard.
4. Hover over the title of the graph, choose **Widget actions, Edit**.
5. On the **Edit graph** screen, in the lower half of the screen, choose the title of the graph.
6. For **Title**, type a new name, choose **Ok** (check mark), and then in the lower-right corner of the **Edit graph** screen, choose **Update widget**.

Add or Remove a Text Widget from a CloudWatch Dashboard

A text widget contains a block of text in [Markdown](#) format. You can add, edit, or remove text widgets from your CloudWatch dashboard.

To add a text widget to a dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose a dashboard, and then choose **Add widget**.
4. In the **Add widget to dashboard** dialog box, choose **Text widget** and then choose **Configure**.
5. In the **New text widget** dialog box, for **Markdown**, add and format your text using [Markdown](#), and then choose **Create widget**.

6. Choose **Save dashboard**.

To edit a text widget on a dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose a dashboard.
4. Hover over the upper-right corner of the text block, and then choose **Widget actions, Edit**.
5. In the **Edit text widget** dialog box, update the text as needed, and then choose **Update widget**.
6. Choose **Save dashboard**.

To remove a text widget from a dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose a dashboard.
4. Hover over the upper-right corner of the text block, and then choose **Widget actions, Delete**.
5. Choose **Save dashboard**.

Monitor Resources in Multiple Regions Using a CloudWatch Dashboard

You can monitor AWS resources in multiple regions using a single CloudWatch dashboard. For example, you can create a dashboard that shows CPU utilization for an EC2 instance located in the `us-west-2` region with your billing metrics, which are located in the `us-east-1` region.

To monitor resources in multiple regions in one dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.
3. In the navigation bar, select a region.
4. Select the metrics to add to your dashboard.
5. For **Actions**, choose **Add to dashboard**.
6. In the **Add to dashboard** dialog box, for **Add to**, type a name for the new dashboard and choose **Add to dashboard**.

Alternatively, to add to an existing dashboard, choose **Existing dashboard**, choose a dashboard, and then choose **Add to dashboard**.

7. Select the next region and repeat these steps to add metrics from this region.
8. Choose **Save dashboard**.

Link and Unlink Graphs on a CloudWatch Dashboard

You can link the graphs on your dashboard together, so that when you zoom in or zoom out on one graph, the other graphs zoom in or zoom out at the same time. You can unlink graphs to limit zoom to one graph.

To link the graphs on a dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose a dashboard.
4. Select **Actions, Link graphs**.

To unlink the graphs on a dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose a dashboard.
4. Clear **Actions, Link graphs**.

Change the Refresh Interval for the CloudWatch Dashboard

You can change how often the data on your CloudWatch dashboard is refreshed or set it to automatically refresh.

To change the dashboard refresh interval

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose a dashboard.
4. On the **Refresh options** menu (upper right corner), choose **1 Minute, 2 Minutes, 5 Minutes, or 15 Minutes**.

To automatically refresh the dashboard

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose a dashboard.
4. On the **Refresh options** menu (upper right corner), select **Auto refresh**.

Change the Time Range or Format of a CloudWatch Dashboard

You can change the time range to display dashboard data over minutes, hours, days, or weeks. You can also change the time format to display dashboard data in UTC or local time.

To change the dashboard time range

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose a dashboard.
4. Do one of the following:

- Select one of the predefined ranges shown, which span from 1 hour to 15 months ago: 1h, 3h, 12h, 1d, 3d, or 1w.
- Choose the **custom** menu and then choose **Relative**. Select one of the predefined ranges, which span from 5 minutes to 15 months.
- Choose the **custom** menu and then choose **Absolute**. Use the calendar picker or the text fields to specify the time range.

To change the dashboard time format

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Dashboards**.
3. Choose a dashboard.
4. Choose the **custom** menu.
5. From the upper corner, choose **UTC** or **Local timezone**.

Using Amazon CloudWatch Metrics

Metrics are data about the performance of your systems. By default, several services provide free metrics for resources (such as Amazon EC2 instances, Amazon EBS volumes, and Amazon RDS DB instances). You can also enable detailed monitoring some resources, such as your Amazon EC2 instances, or publish your own application metrics. Amazon CloudWatch can load all the metrics in your account (both AWS resource metrics and application metrics that you provide) for search, graphing, and alarms.

Metric data is kept for a period of 15 months, enabling you to view both up-to-the-minute data and historical data.

Contents

- [View Available Metrics \(p. 25\)](#)
- [Search for Available Metrics \(p. 28\)](#)
- [Get Statistics for a Metric \(p. 29\)](#)
- [Graph Metrics \(p. 37\)](#)
- [Publish Custom Metrics \(p. 43\)](#)

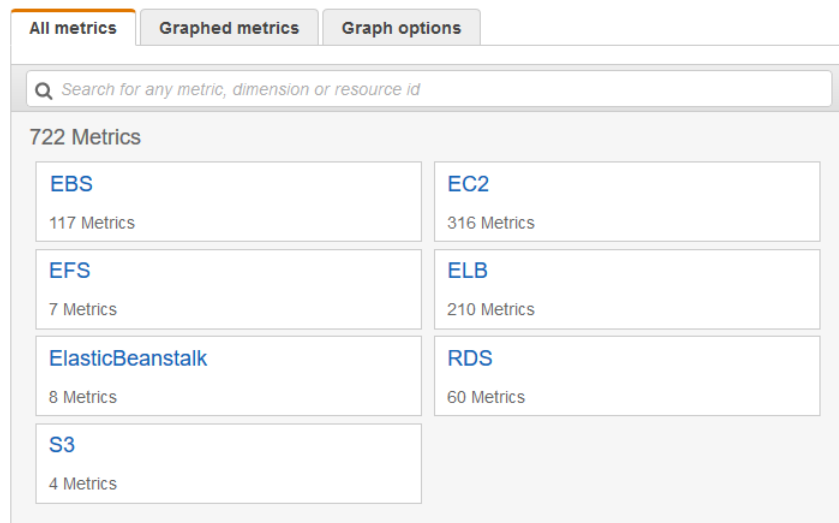
View Available Metrics

Metrics are grouped first by namespace, and then by the various dimension combinations within each namespace. For example, you can view all EC2 metrics, EC2 metrics grouped by instance, or EC2 metrics grouped by Auto Scaling group.

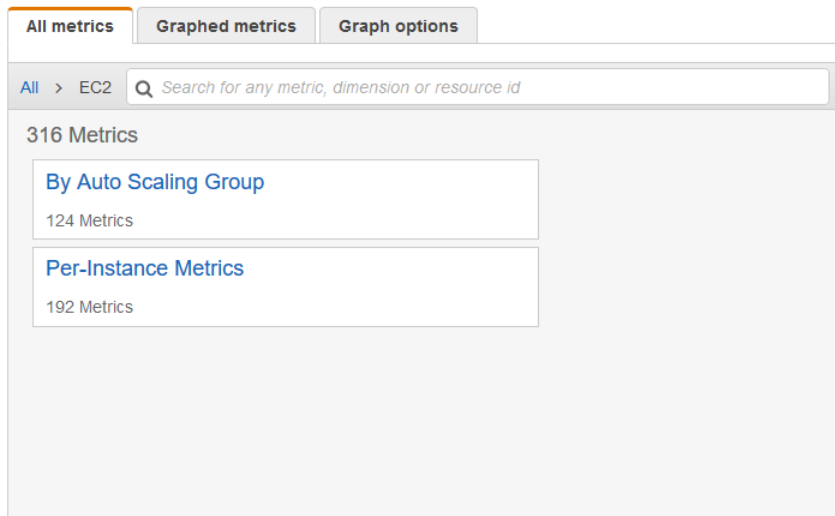
Note that only the AWS services that you're using send metrics to Amazon CloudWatch.

To view available metrics by namespace and dimension using the console

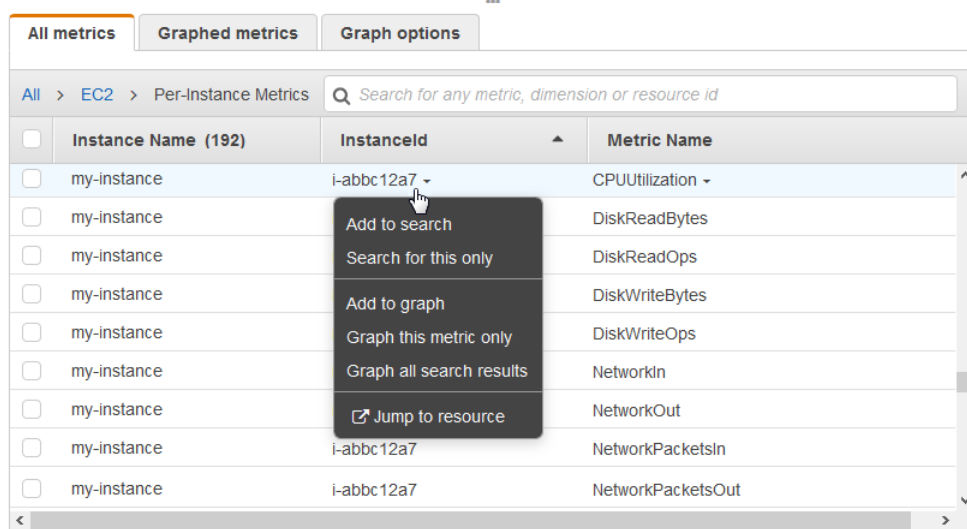
1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.
3. Select a metric namespace (for example, EC2).



4. Select a metric dimension (for example, Per-Instance Metrics).



5. The **All metrics** tab displays all metrics for that dimension in the namespace. You can do the following:
 - a. To sort the table, use the column heading.
 - b. To graph a metric, select the check box next to the metric. To select all metrics, select the check box in the heading row of the table.
 - c. To filter by resource, choose the resource ID and then choose **Add to search**.
 - d. To filter by metric, choose the metric name and then choose **Add to search**.



To view available metrics by namespace, dimension, or metric using the AWS CLI

Use the `list-metrics` command to list CloudWatch metrics. For a list of all service namespaces, see [AWS Namespaces \(p. 47\)](#). For lists of the metrics and dimensions for each service, see [Amazon CloudWatch Metrics and Dimensions Reference \(p. 46\)](#).

The following example specifies the `AWS/EC2` namespace to view all the metrics for Amazon EC2:

```
aws cloudwatch list-metrics --namespace AWS/EC2
```

The following is example output:

```
{
  "Metrics" : [
    ...
    {
      "Namespace": "AWS/EC2",
      "Dimensions": [
        {
          "Name": "InstanceId",
          "Value": "i-1234567890abcdef0"
        }
      ],
      "MetricName": "NetworkOut"
    },
    {
      "Namespace": "AWS/EC2",
      "Dimensions": [
        {
          "Name": "InstanceId",
          "Value": "i-1234567890abcdef0"
        }
      ],
      "MetricName": "CPUUtilization"
    }
  ]
}
```

```
{
  "Namespace": "AWS/EC2",
  "Dimensions": [
    {
      "Name": "InstanceId",
      "Value": "i-1234567890abcdef0"
    }
  ],
  "MetricName": "NetworkIn"
},
...
]
```

To list all the available metrics for a specified resource

The following example specifies the `AWS/EC2` namespace and the `InstanceId` dimension to view the results for the specified instance only.

```
aws cloudwatch list-metrics --namespace AWS/EC2 --dimensions
Name=InstanceId,Value=i-1234567890abcdef0
```

To list a metric for all resources

The following example specifies the `AWS/EC2` namespace and a metric name to view the results for the specified metric only.

```
aws cloudwatch list-metrics --namespace AWS/EC2 --metric-name CPUUtilization
```

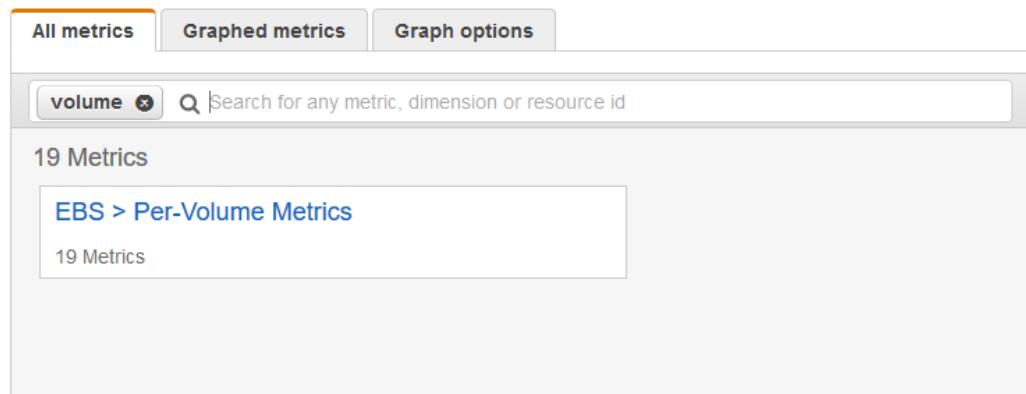
Search for Available Metrics

You can search within all the metrics in your account using targeted search terms. Metrics are returned that have matching results within their namespace, metric name, or dimensions.

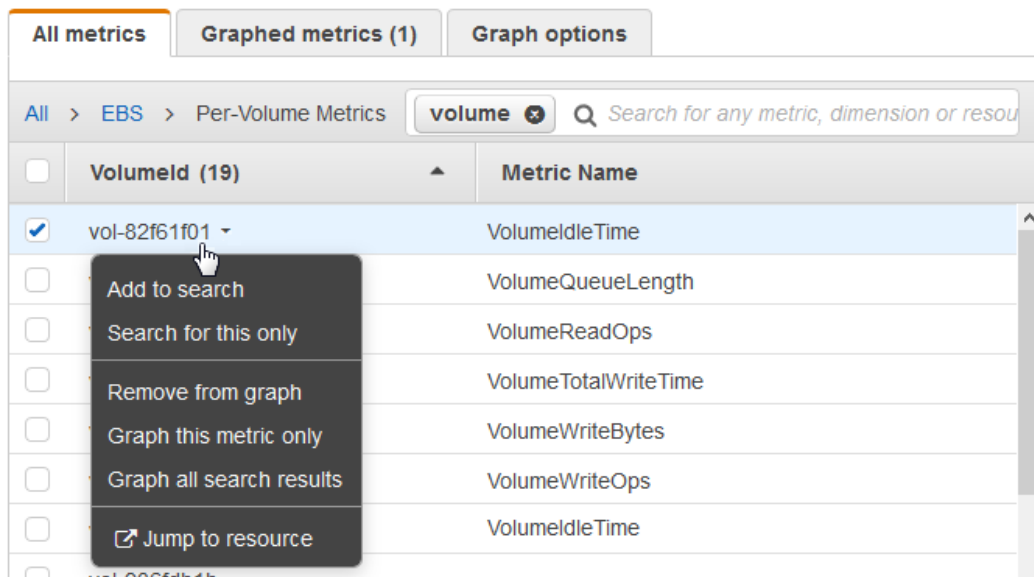
To search for available metrics in CloudWatch

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.
3. In the search field on the **All metrics** tab, type a search term, such as a metric name, service name, or resource name, and press Enter. This shows you all the namespaces with metrics with this search term.

For example, if you search for `volume`, this shows the namespaces that contain metrics with this term in their name.



4. Select a namespace with results for your search to view the metrics. You can do the following:
 - a. To graph one or more metrics, select the check box next to each metric. To select all metrics, select the check box in the heading row of the table.
 - b. To view one of the resources in its console, choose the resource ID and then choose **Jump to resource**.
 - c. To view help for a metric, choose the metric name and then choose **What is this?**



Get Statistics for a Metric

The following examples show you how to get statistics for the CloudWatch metrics for your resources, such as your EC2 instances.

Examples

- [Get Statistics for a Specific Resource \(p. 30\)](#)
- [Aggregate Statistics Across Resources \(p. 33\)](#)
- [Aggregate Statistics by Auto Scaling Group \(p. 34\)](#)
- [Aggregate Statistics by Amazon Machine Image \(AMI\) \(p. 36\)](#)

Get Statistics for a Specific Resource


The following example shows you how to determine the maximum CPU utilization of a specific EC2 instance.

Requirements

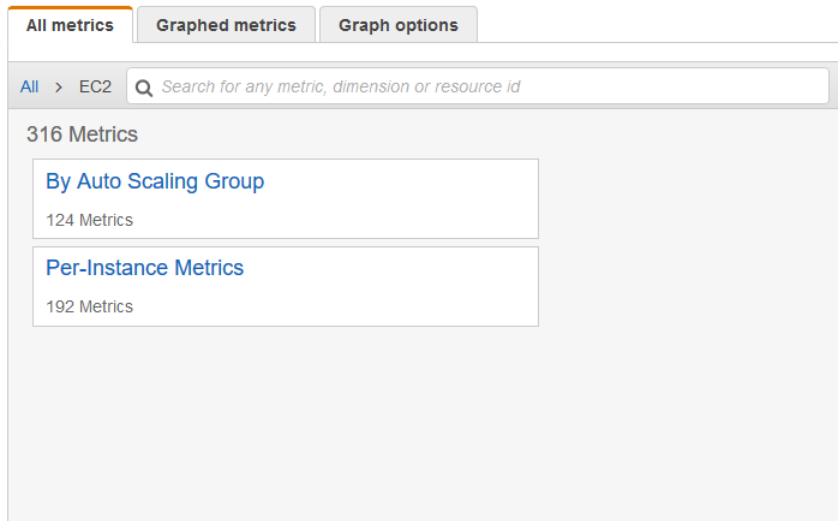
- You must have the ID of the instance. You can get the instance ID using the Amazon EC2 console or the `describe-instances` command.
- By default, basic monitoring is enabled, but you can enable detailed monitoring. For more information, see [Enable or Disable Detailed Monitoring for Your Instances](#) in the *Amazon EC2 User Guide for Linux Instances*.

To display the average CPU utilization for a specific instance using the console

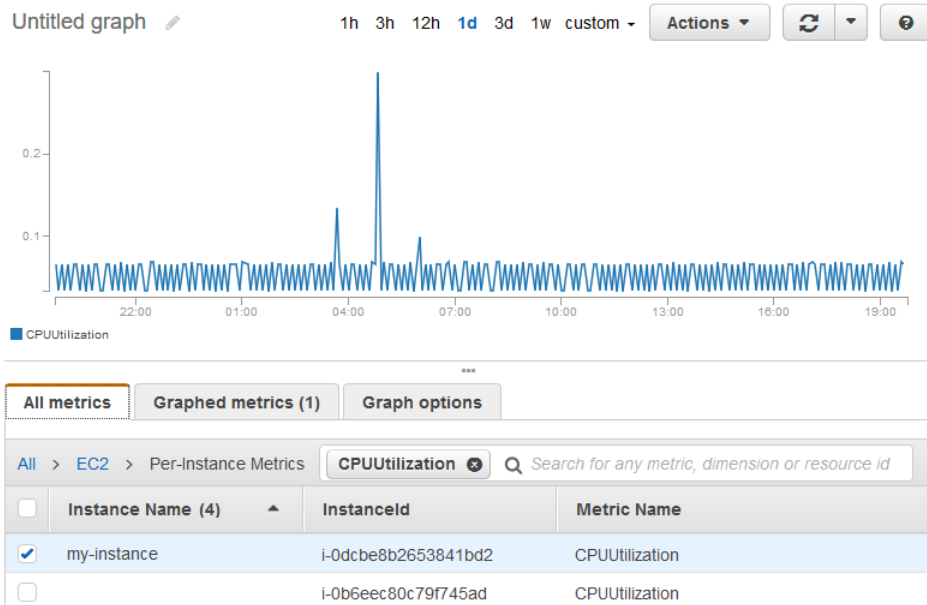
1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.
3. Select the EC2 metric namespace.



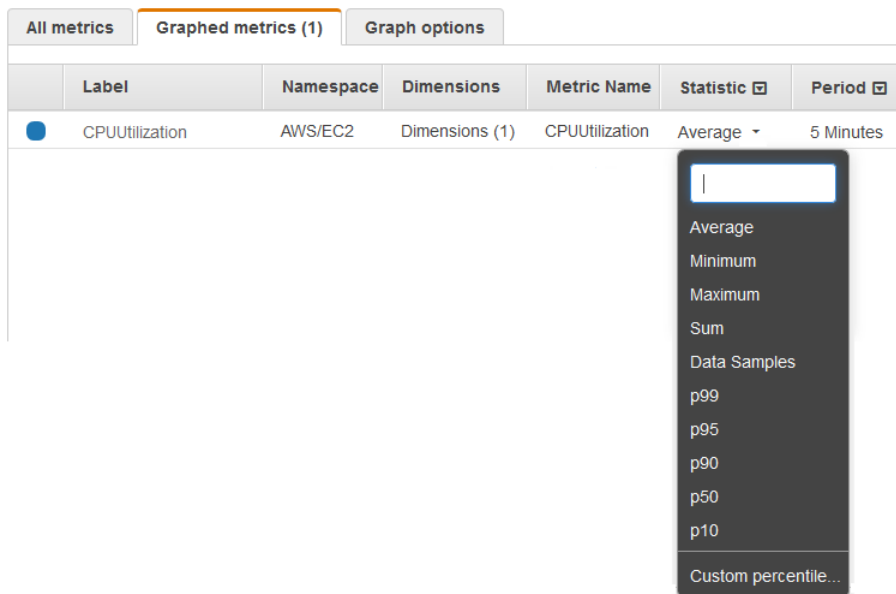
4. Select the Per-Instance Metrics dimension.



- In the search field, type `CPUUtilization` and press Enter. Select the row for the specific instance, which displays a graph for the **CPUUtilization** metric for the instance. To change the name of the graph, choose the pencil icon. To change the time range, select one of the predefined values or choose **custom**.



- To change the statistic, choose the **Graphed metrics** tab. Choose the column heading or an individual value, and then choose one of the statistics or predefined percentiles, or specify a custom percentile (for example, p95.45).



- To change the period, choose the **Graphed metrics** tab. Choose the column heading or an individual value, and then choose a different value.

To get the CPU utilization per EC2 instance using the AWS CLI

Use the [get-metric-statistics](#) command as follows to get the **CPUUtilization** metric for the specified instance:

```
aws cloudwatch get-metric-statistics --namespace AWS/EC2 --metric-name
CPUUtilization \
--dimensions Name=InstanceId,Value=i-1234567890abcdef0 --statistics Maximum \
--start-time 2016-10-18T23:18:00 --end-time 2016-10-19T23:18:00 --period 360
```

The returned statistics are six-minute values for the requested two-day time interval. Each value represents the maximum CPU utilization percentage for a single EC2 instance. The following is example output:

```
{
  "Datapoints": [
    {
      "Timestamp": "2016-10-19T00:18:00Z",
      "Maximum": 0.33000000000000002,
      "Unit": "Percent"
    },
    {
      "Timestamp": "2016-10-19T03:18:00Z",
      "Maximum": 99.670000000000002,
      "Unit": "Percent"
    },
    {
      "Timestamp": "2016-10-19T07:18:00Z",
      "Maximum": 0.34000000000000002,
      "Unit": "Percent"
    }
  ]
}
```

```
    ...  
  ],  
  "Label": "CPUUtilization"  
}
```

Aggregate Statistics Across Resources

You can aggregate the metrics for AWS resources across multiple resources. Note that Amazon CloudWatch cannot aggregate data across regions. Metrics are completely separate between regions.

For example, you can aggregate statistics for your EC2 instances that have detailed monitoring enabled. Instances that use basic monitoring are not included. Therefore, you must enable detailed monitoring (at an additional charge), which provides data in 1-minute periods. For more information, see [Enable or Disable Detailed Monitoring for Your Instances](#) in the *Amazon EC2 User Guide for Linux Instances*.

This example shows you how to get the average CPU usage for your EC2 instances. Because no dimension is specified, CloudWatch returns statistics for all dimensions in the `AWS/EC2` namespace. To get statistics for other metrics, see [Amazon CloudWatch Metrics and Dimensions Reference](#) (p. 46).


Important

This technique for retrieving all dimensions across an AWS namespace does not work for custom namespaces that you publish to Amazon CloudWatch. With custom namespaces, you must specify the complete set of dimensions that are associated with any given data point to retrieve statistics that include the data point.

To display average CPU utilization for your EC2 instances

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.
3. Select the **EC2** namespace and then select **Across All Instances**.
4. Select the row that contains **CPUUtilization**, which displays a graph for the metric for all your EC2 instances. To change the name of the graph, choose the pencil icon. To change the time range, select one of the predefined values or choose **custom**.

Untitled graph  1h 3h 12h 1d 3d 1w custom   



All metrics	Graphed metrics (1)	Graph options
All > EC2 > Across All Instances	<input type="text" value="Search for any metric, dimension or resource id"/>	
<input type="checkbox"/>	Metric Name (7)	
<input checked="" type="checkbox"/>	CPUUtilization	
<input type="checkbox"/>	DiskReadBytes	
<input type="checkbox"/>	DiskReadOps	

5. To change the statistic, choose the **Graphed metrics** tab. Choose the column heading or an individual value, and then choose one of the statistics or predefined percentiles, or specify a custom percentile (for example, p95.45).
6. To change the period, choose the **Graphed metrics** tab. Choose the column heading or an individual value, and then choose a different value.

To get average CPU utilization across your EC2 instances using the AWS CLI

Use the `get-metric-statistics` command as follows:

```
aws cloudwatch get-metric-statistics --namespace AWS/EC2 --metric-name  
CPUUtilization --statistics "Average" "SampleCount" \  
--start-time 2016-10-11T23:18:00 --end-time 2016-10-12T23:18:00 --period 3600
```

The following is example output:

```
{  
  "Datapoints": [  
    {  
      "SampleCount": 238.0,  
      "Timestamp": "2016-10-12T07:18:00Z",  
      "Average": 0.038235294117647062,  
      "Unit": "Percent"  
    },  
    {  
      "SampleCount": 240.0,  
      "Timestamp": "2016-10-12T09:18:00Z",  
      "Average": 0.16670833333333332,  
      "Unit": "Percent"  
    },  
    {  
      "SampleCount": 238.0,  
      "Timestamp": "2016-10-11T23:18:00Z",  
      "Average": 0.041596638655462197,  
      "Unit": "Percent"  
    },  
    ...  
  ],  
  "Label": "CPUUtilization"  
}
```

Aggregate Statistics by Auto Scaling Group

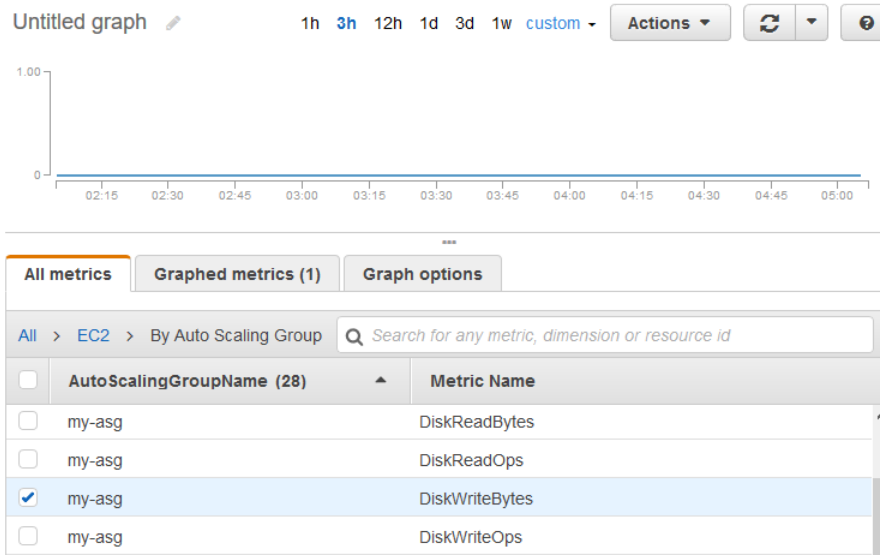
You can aggregate statistics for the EC2 instances in an Auto Scaling group. Note that Amazon CloudWatch cannot aggregate data across regions. Metrics are completely separate between regions.

This example shows you how to get the total bytes written to disk for one Auto Scaling group. The total is computed for one-minute periods for a 24-hour interval across all EC2 instances in the specified Auto Scaling group.

To display DiskWriteBytes for the instances in an Auto Scaling group using the console

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.
3. Select the **EC2** namespace and then select **By Auto Scaling Group**.

- Select the row for the **DiskWriteBytes** metric and the specific Auto Scaling group, which displays a graph for the metric for the instances in the Auto Scaling group. To change the name of the graph, choose the pencil icon. To change the time range, select one of the predefined values or choose **custom**.



- To change the statistic, choose the **Graphed metrics** tab. Choose the column heading or an individual value, and then choose one of the statistics or predefined percentiles, or specify a custom percentile (for example, p95.45).
- To change the period, choose the **Graphed metrics** tab. Choose the column heading or an individual value, and then choose a different value.

To get DiskWriteBytes for the instances in an Auto Scaling group using the AWS CLI

Use the `get-metric-statistics` command as follows:

```
aws cloudwatch get-metric-statistics --namespace AWS/EC2 --metric-name
DiskWriteBytes
--dimensions Name=AutoScalingGroupName,Value=my-asg --statistics "Sum"
"SampleCount" \
--start-time 2016-10-16T23:18:00 --end-time 2016-10-18T23:18:00 --period 360
```

The following is example output:

```
{
  "Datapoints": [
    {
      "SampleCount": 18.0,
      "Timestamp": "2016-10-19T21:36:00Z",
      "Sum": 0.0,
      "Unit": "Bytes"
    },
    {
      "SampleCount": 5.0,
      "Timestamp": "2016-10-19T21:42:00Z",
      "Sum": 0.0,
      "Unit": "Bytes"
    }
  ]
}
```

```
],  
  "Label": "DiskWriteBytes"  
}
```

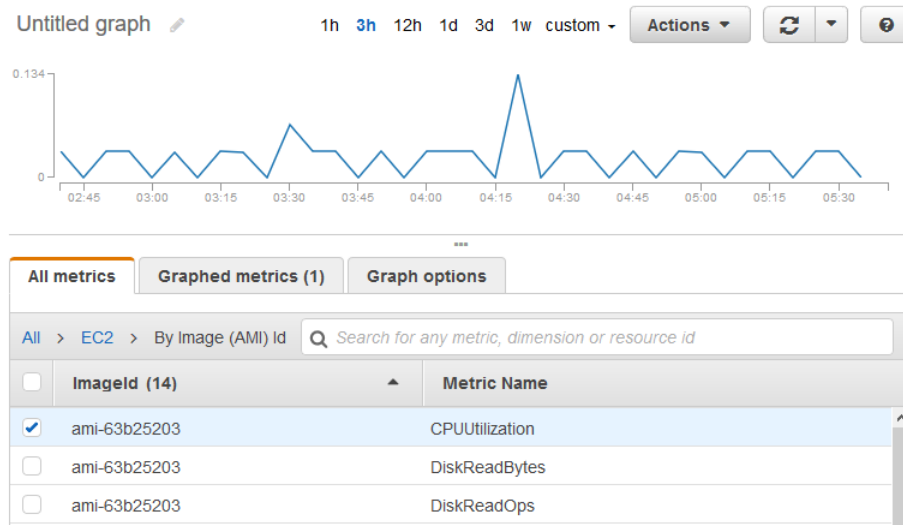
Aggregate Statistics by Amazon Machine Image (AMI)

You can aggregate statistics for the EC2 instances that have detailed monitoring enabled. Instances that use basic monitoring are not included. For more information, see [Enable or Disable Detailed Monitoring for Your Instances](#) in the *Amazon EC2 User Guide for Linux Instances*.

This example shows you how to determine average CPU utilization for all instances that use the specified AMI. The average is over 60-second time intervals for a one-day period.

To display the average CPU utilization by AMI using the console

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.
3. Select the **EC2** namespace and then select **By Image (AMI) Id**.
4. Select the row for the **CPUUtilization** metric and the specific AMI, which displays a graph for the metric for the specified AMI. To change the name of the graph, choose the pencil icon. To change the time range, select one of the predefined values or choose **custom**.



5. To change the statistic, choose the **Graphed metrics** tab. Choose the column heading or an individual value, and then choose one of the statistics or predefined percentiles, or specify a custom percentile (for example, p95.45).
6. To change the period, choose the **Graphed metrics** tab. Choose the column heading or an individual value, and then choose a different value.

To get the average CPU utilization by AMI using the AWS CLI

Use the `get-metric-statistics` command as follows:

```
aws cloudwatch get-metric-statistics --namespace AWS/EC2 --metric-name  
CPUUtilization \
```

```
--dimensions Name=ImageId,Value=ami-3c47a355 --statistics Average \  
--start-time 2016-10-10T00:00:00 --end-time 2016-10-11T00:00:00 --period 3600
```

The operation returns statistics that are one-minute values for the one-day interval. Each value represents an average CPU utilization percentage for EC2 instances running the specified AMI. The following is example output:

```
{  
  "Datapoints": [  
    {  
      "Timestamp": "2016-10-10T07:00:00Z",  
      "Average": 0.041000000000000009,  
      "Unit": "Percent"  
    },  
    {  
      "Timestamp": "2016-10-10T14:00:00Z",  
      "Average": 0.079579831932773085,  
      "Unit": "Percent"  
    },  
    {  
      "Timestamp": "2016-10-10T06:00:00Z",  
      "Average": 0.0360000000000000011,  
      "Unit": "Percent"  
    },  
    ...  
  ],  
  "Label": "CPUUtilization"  
}
```

Graph Metrics

You can use the CloudWatch console to graph metric data generated by other AWS services to make it easier to see the metric activity on your services. You can use the following procedures to graph metrics in CloudWatch.

Contents

- [Graph a Metric \(p. 37\)](#)
- [Modify the Time Range for a Graph \(p. 40\)](#)
- [Modify the Y Axis for a Graph \(p. 42\)](#)
- [Create an Alarm from a Metric on a Graph \(p. 42\)](#)

Graph a Metric

You can select metrics and create graphs of the data using the CloudWatch console.

CloudWatch supports the following statistics on metrics: Average, Minimum, Maximum, Sum, and SampleCount. For more information, see [Statistics \(p. 5\)](#).

You can view your data at different granularities. For example, you can choose a detailed view (for example 1 minute), which can be useful when troubleshooting. You can choose a less detailed view (for example, 1 hour), which can be useful when viewing a broader time range (for example, 3 days) so that you can see trends over time. For more information, see [Periods \(p. 6\)](#).

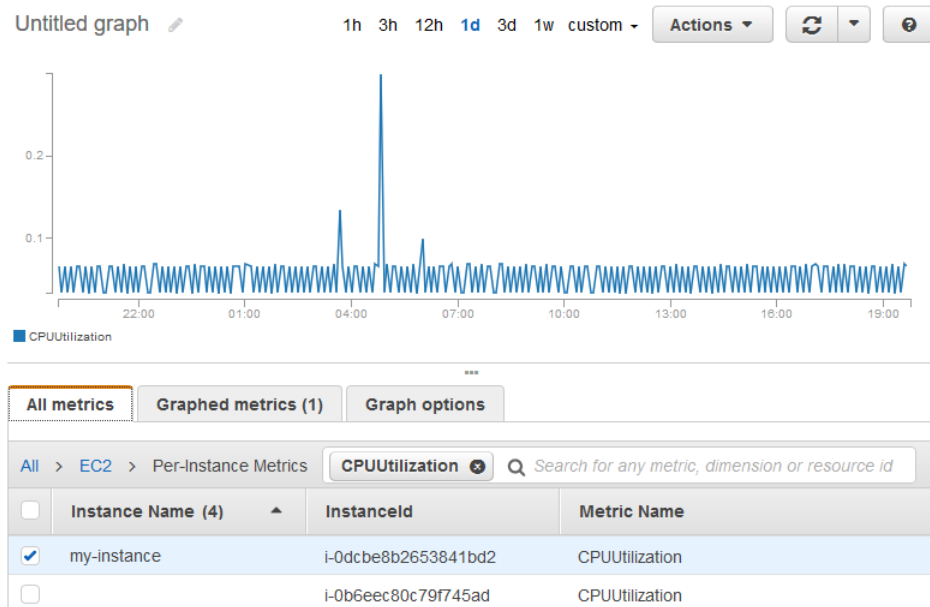
Create a Graph

To graph a metric

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.
3. On the **All metrics** tab, type a search term in the search field, such as a metric name or resource name, and press Enter.

For example, if you search for the **CPUUtilization** metric, you'll see the namespaces and dimensions with this metric.

4. Select one of the results for your search to view the metrics.
5. To graph one or more metrics, select the check box next to each metric. To select all metrics, select the check box in the heading row of the table.

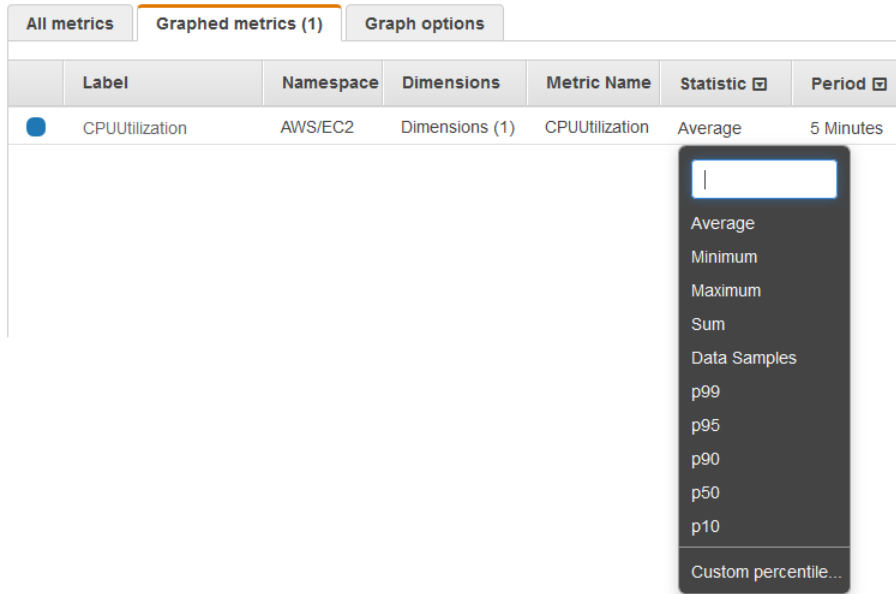


6. To view more information about the metric being graphed, hover over the legend.
7. To get a URL for your graph, choose **Actions**, **Share**. Copy the URL and save it or share it.
8. To add your graph to a dashboard, choose **Actions**, **Add to dashboard**.

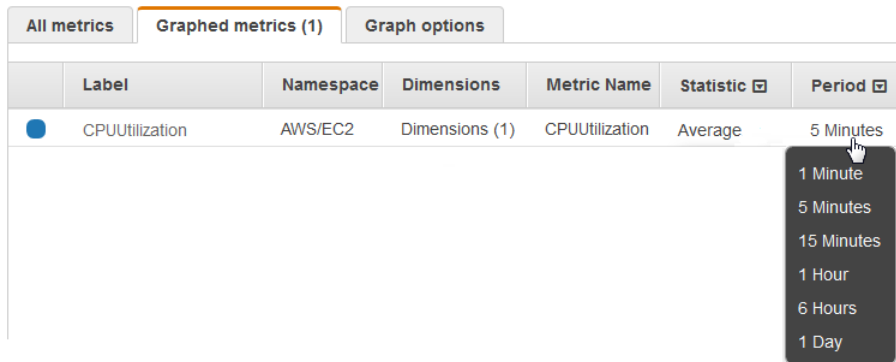
Update a Graph

To update your graph

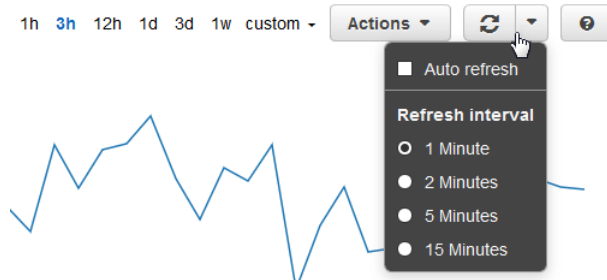
1. To change the name of the graph, choose the pencil icon.
2. To change the time range, select one of the predefined values or choose **custom**. For more information, see [Modify the Time Range for a Graph \(p. 40\)](#).
3. To change the statistic, choose the **Graphed metrics** tab. Choose the column heading or an individual value, and then choose one of the statistics or predefined percentiles, or specify a custom percentile (for example, p95.45).



- To change the period, choose the **Graphed metrics** tab. Choose the column heading or an individual value, and then choose a different value.




- To change the refresh interval, choose **Refresh options**, and then select **Auto refresh** or choose **1 Minute**, **2 Minutes**, **5 Minutes**, or **15 Minutes**.



Duplicate a Metric

To duplicate a metric

- Choose the **Graphed metrics** tab.
- For **Actions**, choose the **Duplicate** icon.



3. Update the duplicate metric as needed.

Modify the Time Range for a Graph


You can change the time range for a graph to view the data at different points in time.

Relative Time Ranges

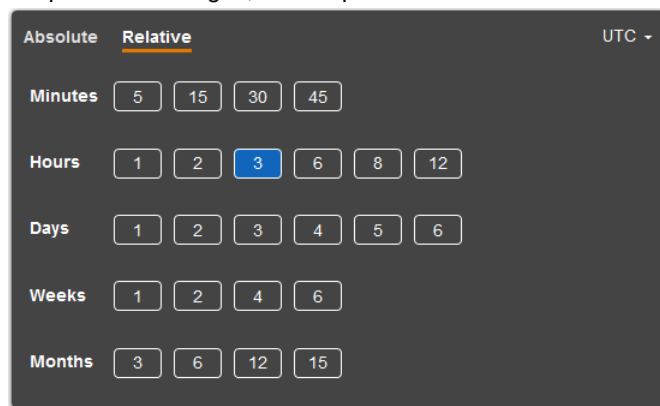
You can set a relative time range for your graph.

To specify a relative time range for a graph

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.
3. Select one of the predefined ranges shown at the top of the page, which span from 1 hour to 1 week ago.



4. For more predefined ranges, choose the **custom** menu and then choose **Relative**. Select one of the predefined ranges, which span from 5 minutes to 15 months ago.

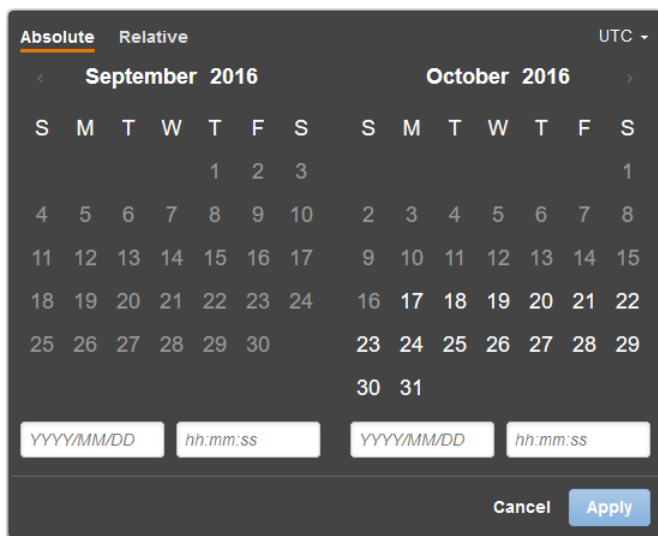


Absolute Time Ranges

You can set an absolute time range for your graph.

To specify an absolute time range for a graph

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.
3. Choose the **custom** menu and then choose **Absolute**. Use the calendar picker or the text fields to specify the time range.



Zoom in on a Graph


You can change the granularity of a graph and zoom in to see data over a shorter time period.

To zoom in on a graph

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.
3. Choose and drag on the graph area, and then release your mouse button.



4. To reset a zoomed-in graph, choose the **Reset zoom** icon.



Modify the Y Axis for a Graph

You can set custom bounds for the Y axis on a graph to help you see the data better. For example, you can change the bounds on a CPUUtilization graph to 100 percent so that it's easy to see whether the CPU is low (the plotted line is near the bottom of the graph) or high (the plotted line is near the top of the graph).

You can switch between two different Y axes for your graph. This is particularly useful if the graph contains metrics that have different units or that differ greatly in their range of values.

To modify the Y axis on a graph

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.
3. Select a metric namespace (for example, EC2) and then a metric dimension (for example, Per-Instance Metrics).
4. The **All metrics** tab displays all metrics for that dimension in that namespace. To graph a metric, select the check box next to the metric.
5. On the **Graph options** tab, specify the **Min** and **Max** values for **Left Y Axis**. Note that the value of **Min** cannot be greater than the value of **Max**.

The screenshot shows the 'Graph options' tab in the CloudWatch console. It has three sub-tabs: 'All metrics', 'Graphed metrics (1)', and 'Graph options'. The 'Graph options' tab is active. Under 'Left Y Axis', there is a 'Limits' section with 'Min' set to 0 and 'Max' set to 100. Under 'Right Y Axis', there is a 'Limits' section with 'Min' set to Auto and 'Max' set to Auto.

6. To create a second Y axis, specify the **Min** and **Max** values for **Right Y Axis**.
7. To switch between the two Y axes, choose the **Graphed metrics** tab. For **Y Axis**, choose **Left Y Axis** or **Right Y Axis**.

The screenshot shows the 'Graphed metrics (1)' tab in the CloudWatch console. The 'Graph options' sub-tab is active. Below the sub-tabs is a table with columns: Namespace, Dimensions, Metric Name, Statistic, Period, Y Axis, and Actions. The table contains one row: 'lization', 'AWS/EC2', 'Dimensions (1)', 'CPUUtilization', 'Average', '5 Minutes', and 'Y Axis'. The 'Y Axis' dropdown menu is open, showing 'Right Y Axis' as the selected option.

Namespace	Dimensions	Metric Name	Statistic	Period	Y Axis	Actions
lization	AWS/EC2	Dimensions (1)	CPUUtilization	Average	5 Minutes	Y Axis


Create an Alarm from a Metric on a Graph

You can graph a metric and then create an alarm from the metric on the graph, which has the benefit of populating many of the alarm fields for you.

To create an alarm from a metric on a graph

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Metrics**.

3. Select a metric namespace (for example, EC2) and then a metric dimension (for example, Per-Instance Metrics).
4. The **All metrics** tab displays all metrics for that dimension in that namespace. To graph a metric, select the check box next to the metric.
5. To create an alarm for the metric, choose the **Graphed metrics** tab. For **Actions**, choose the alarm icon.



6. Under **Alarm Threshold**, type a unique name for the alarm and a description of the alarm. For **Whenever**, specify a threshold and the number of periods.
7. Under **Actions**, select the type of action you want the alarm to perform when the alarm is triggered.
8. (Optional) For **Period**, choose a different value. For **Statistic**, choose **Standard** to specify one of the statistics in the list or choose **Custom** to specify a percentile (for example, p95.45).

Period:

Statistic: Standard Custom

9. Choose **Create Alarm**.

Publish Custom Metrics

You can publish your own metrics to CloudWatch using the AWS CLI or an API. You can view statistical graphs of your published metrics with the AWS Management Console.

CloudWatch stores data about a metric as a series of data points. Each data point has an associated time stamp. You can even publish an aggregated set of data points called a *statistics set*.

Topics

- [Publish Single Data Points \(p. 43\)](#)
- [Publish Statistic Sets \(p. 44\)](#)
- [Publish the Value Zero \(p. 45\)](#)

Publish Single Data Points

To publish a single data point for a new or existing metric, use the `put-metric-data` command with one value and time stamp. For example, the following actions each publish one data point:

```
aws cloudwatch put-metric-data --metric-name PageViewCount --
namespace MyService --value 2 --timestamp 2016-10-14T12:00:00.000Z
aws cloudwatch put-metric-data --metric-name PageViewCount --
namespace MyService --value 4 --timestamp 2016-10-14T12:00:01.000Z
aws cloudwatch put-metric-data --metric-name PageViewCount --
namespace MyService --value 5 --timestamp 2016-10-14T12:00:02.000Z
```

If you call this command with a new metric name, CloudWatch creates a metric for you. Otherwise, CloudWatch associates your data with the existing metric that you specified.

Note

When you create a metric, it can take up to two minutes before you can retrieve statistics for the new metric using the [get-metric-statistics](#) command. However, it can take up to fifteen minutes before the new metric appears in the list of metrics retrieved using the [list-metrics](#) command.

Although you can publish data points with time stamps as granular as one-thousandth of a second, CloudWatch aggregates the data to a minimum granularity of one minute. CloudWatch records the average (sum of all items divided by number of items) of the values received for every 1-minute period, as well as number of samples, maximum value, and minimum value for the same time period. For example, the `PageViewCount` metric from the previous examples contains three data points with time stamps just seconds apart. CloudWatch aggregates the three data points because they all have time stamps within a one-minute period.

CloudWatch uses one-minute boundaries when aggregating data points. For example, CloudWatch aggregates the data points from the previous example because all three data points fall within the one-minute period that begins at `2016-10-20T12:00:00.000Z` and ends at `2016-10-20T12:01:00.000Z`.

You can use the [get-metric-statistics](#) command to retrieve statistics based on the data points that you published.

```
aws cloudwatch get-metric-statistics--namespace MyService --metric-  
name PageViewCount \  
--statistics "Sum" "Maximum" "Minimum" "Average" "SampleCount" \  
--start-time 2016-10-20T12:00:00.000Z --end-time 2016-10-20T12:05:00.000Z --  
period 60
```

The following is example output:

```
{  
  "Datapoints": [  
    {  
      "SampleCount": 3.0,  
      "Timestamp": "2016-10-20T12:00:00Z",  
      "Average": 3.6666666666666665,  
      "Maximum": 5.0,  
      "Minimum": 2.0,  
      "Sum": 11.0,  
      "Unit": "None"  
    }  
  ],  
  "Label": "PageViewCount"  
}
```

Publish Statistic Sets

You can aggregate your data before you publish to CloudWatch. When you have multiple data points per minute, aggregating data minimizes the number of calls to [put-metric-data](#). For example, instead of calling [put-metric-data](#) multiple times for three data points that are within three seconds of each other, you can aggregate the data into a statistic set that you publish with one call:

```
aws cloudwatch put-metric-data --metric-name  
PageViewCount --namespace MyService --statistic-
```

```
value Sum=11,Minimum=2,Maximum=5,SampleCount=3 --  
timestamp 2016-10-14T12:00:00.000Z
```

Publish the Value Zero

When your data is more sporadic and you have periods that have no associated data, you can choose to publish the value zero (0) for that period or no value at all. You might want to publish zero instead of no value if you use periodic calls to `PutMetricData` to monitor the health of your application. For example, you can set a CloudWatch alarm to notify you if your application fails to publish metrics every five minutes. You want such an application to publish zeros for periods with no associated data.

You might also publish zeros if you want to track the total number of data points or if you want statistics such as minimum and average to include data points with the value 0.

Amazon CloudWatch Metrics and Dimensions Reference

This reference includes all the namespaces, dimensions, and metrics that you can use with CloudWatch. Namespaces are containers for metrics. Metrics, which are time-ordered sets of data points, are isolated from one another in different namespaces so that metrics from different applications are not mistakenly aggregated into the same statistics. In addition, each metric has a dimension, which is a name/value pair that you can use to filter metrics.

Metrics and Dimensions

- [AWS Namespaces \(p. 47\)](#)
- [API Gateway \(p. 48\)](#)
- [Auto Scaling \(p. 50\)](#)
- [AWS Billing and Cost Management \(p. 51\)](#)
- [Amazon CloudFront \(p. 51\)](#)
- [Amazon CloudSearch \(p. 53\)](#)
- [Amazon CloudWatch Events \(p. 54\)](#)
- [Amazon CloudWatch Logs \(p. 55\)](#)
- [Amazon DynamoDB \(p. 56\)](#)
- [Amazon EC2 \(p. 66\)](#)
- [Amazon EC2 Spot Fleet \(p. 70\)](#)
- [Amazon ECS \(p. 71\)](#)
- [Elastic Beanstalk \(p. 73\)](#)
- [Amazon ElastiCache \(p. 74\)](#)
- [Amazon EBS \(p. 80\)](#)
- [Amazon EFS \(p. 82\)](#)
- [Elastic Load Balancing \(p. 84\)](#)
- [Amazon EMR \(p. 89\)](#)
- [Amazon ES \(p. 99\)](#)
- [Elastic Transcoder \(p. 101\)](#)
- [AWS IoT \(p. 103\)](#)
- [Amazon Kinesis Analytics \(p. 105\)](#)

- [Amazon Kinesis Firehose](#) (p. 105)
- [Amazon Kinesis Streams](#) (p. 108)
- [AWS KMS](#) (p. 114)
- [AWS Lambda](#) (p. 114)
- [Amazon Machine Learning](#) (p. 116)
- [AWS OpsWorks](#) (p. 117)
- [Amazon Polly](#) (p. 118)
- [Amazon Redshift](#) (p. 119)
- [Amazon RDS](#) (p. 122)
- [Amazon Route 53](#) (p. 124)
- [Amazon SES](#) (p. 126)
- [Amazon SNS](#) (p. 126)
- [Amazon SQS](#) (p. 128)
- [Amazon S3](#) (p. 130)
- [Amazon SWF](#) (p. 133)
- [AWS Storage Gateway](#) (p. 134)
- [AWS WAF](#) (p. 145)
- [Amazon WorkSpaces](#) (p. 146)

AWS Namespaces

CloudWatch namespaces are containers for metrics. Metrics in different namespaces are isolated from each other, so that metrics from different applications are not mistakenly aggregated into the same statistics. All AWS services that provide Amazon CloudWatch data use a namespace string, beginning with "AWS/". When you create custom metrics, you must also specify a namespace as a container for custom metrics. The following services push metric data points to CloudWatch.

AWS Product	Namespace
Amazon API Gateway	AWS/ApiGateway
Auto Scaling	AWS/AutoScaling
AWS Billing	AWS/Billing
Amazon CloudFront	AWS/CloudFront
Amazon CloudSearch	AWS/CloudSearch
Amazon CloudWatch Events	AWS/Events
Amazon CloudWatch Logs	AWS/Logs
Amazon DynamoDB	AWS/DynamoDB
Amazon EC2	AWS/EC2
Amazon EC2	AWS/EC2Spot (Spot Instances)
Amazon EC2 Container Service	AWS/ECS
AWS Elastic Beanstalk	AWS/ElasticBeanstalk
Amazon Elastic Block Store	AWS/EBS

AWS Product	Namespace
Amazon Elastic File System	AWS/EFS
Elastic Load Balancing	AWS/ELB (Classic Load Balancers)
Elastic Load Balancing	AWS/ApplicationELB (Application Load Balancers)
Amazon Elastic Transcoder	AWS/ElasticTranscoder
Amazon ElastiCache	AWS/ElastiCache
Amazon Elasticsearch Service	AWS/ES
Amazon EMR	AWS/ElasticMapReduce
AWS IoT	AWS/IoT
AWS Key Management Service	AWS/KMS
Amazon Kinesis Analytics	AWS/KinesisAnalytics
Amazon Kinesis Firehose	AWS/Firehose
Amazon Kinesis Streams	AWS/Kinesis
AWS Lambda	AWS/Lambda
Amazon Machine Learning	AWS/ML
AWS OpsWorks	AWS/OpsWorks
Amazon Polly	AWS/Polly
Amazon Redshift	AWS/Redshift
Amazon Relational Database Service	AWS/RDS
Amazon Route 53	AWS/Route53
Amazon Simple Email Service	AWS/SES
Amazon Simple Notification Service	AWS/SNS
Amazon Simple Queue Service	AWS/SQS
Amazon Simple Storage Service	AWS/S3
Amazon Simple Workflow Service	AWS/SWF
AWS Storage Gateway	AWS/StorageGateway
AWS WAF	AWS/WAF
Amazon WorkSpaces	AWS/WorkSpaces

Amazon API Gateway Metrics and Dimensions

The metrics and dimensions that API Gateway sends to Amazon CloudWatch are listed below. For more information, see [Monitor API Execution with Amazon CloudWatch](#) in the *Amazon API Gateway Developer Guide*.

API Gateway Metrics

Amazon API Gateway sends metric data to CloudWatch every minute.

The `AWS/ApiGateway` namespace includes the following metrics.

Metric	Description
4XXError	The number of client-side errors captured Unit: count
5XXError	The number of server-side errors captured. Unit: count
CacheHitCount	The number of requests served from the API cache. Unit: count
CacheMissCount	The number of requests served from the back end when API caching is enabled. Unit: count
Count	The number of calls to API methods. Unit: count
IntegrationLatency	The time between when API Gateway relays a request to the back end and when it receives a response from the back end. Unit: millisecond
Latency	The time between when API Gateway receives a request from a client and when it returns a response to the client. Unit: millisecond

Dimensions for Metrics

You can use the dimensions in the following table to filter API Gateway metrics.

Dimension	Description
ApiName	Filters API Gateway metrics for an API of the specified API name.
ApiName, Method, Resource, Stage	Filters API Gateway metrics for an API method of the specified API, stage, resource, and method. API Gateway will not send such metrics unless you have explicitly enabled detailed CloudWatch metrics. You can do this in the console by selecting Enable CloudWatch Metrics under a stage Settings tab. Alternatively, you can call the stage:update

Dimension	Description
	<p>action of the API Gateway REST API to update the <code>metricsEnabled</code> property to <code>true</code>.</p> <p>Enabling such metrics will incur additional charges to your account. For pricing information, see Amazon CloudWatch Pricing.</p>
ApiName, Stage	Filters API Gateway metrics for an API stage of the specified API and stage.

Auto Scaling Metrics and Dimensions

Auto Scaling sends metrics for instances and groups to CloudWatch. For Auto Scaling instances, you can enable detailed (one-minute) monitoring or basic (five-minute) monitoring. For Auto Scaling groups, you can enable group metrics. For more information, see [Monitoring Your Auto Scaling Instances and Groups](#) in the *Auto Scaling User Guide*.

Auto Scaling Group Metrics

If you enable group metrics, Auto Scaling sends aggregated data to CloudWatch every minute.

The `AWS/AutoScaling` namespace includes the following metrics.

Metric	Description
<code>GroupMinSize</code>	The minimum size of the Auto Scaling group.
<code>GroupMaxSize</code>	The maximum size of the Auto Scaling group.
<code>GroupDesiredCapacity</code>	The number of instances that the Auto Scaling group attempts to maintain.
<code>GroupInServiceInstances</code>	The number of instances that are running as part of the Auto Scaling group. This metric does not include instances that are pending or terminating.
<code>GroupPendingInstances</code>	The number of instances that are pending. A pending instance is not yet in service. This metric does not include instances that are in service or terminating.
<code>GroupStandbyInstances</code>	The number of instances that are in a <code>Standby</code> state. Instances in this state are still running but are not actively in service.
<code>GroupTerminatingInstances</code>	The number of instances that are in the process of terminating. This metric does not include instances that are in service or pending.
<code>GroupTotalInstances</code>	The total number of instances in the Auto Scaling group. This metric identifies the number of instances that are in service, pending, and terminating.

Dimensions for Auto Scaling Group Metrics

To filter the metrics for your Auto Scaling group by group name, use the `AutoScalingGroupName` dimension.

AWS Billing and Cost Management Dimensions and Metrics

The AWS Billing and Cost Management service sends metrics to CloudWatch. For more information, see [Monitoring Charges with Alerts and Notifications](#) in the *AWS Billing and Cost Management User Guide*.

AWS Billing and Cost Management Metrics

The `AWS/Billing` namespace includes the following metrics.

Metric	Description
<code>EstimatedCharges</code>	The estimated charges for your AWS usage. This can either be estimated charges for one service or a roll-up of estimated charges for all services.

Dimensions for AWS Billing and Cost Management Metrics

Billing and Cost Management supports filtering metrics by the following dimensions.

Dimension	Description
<code>ServiceName</code>	The name of the AWS service. This dimension is omitted for the total of estimated charges across all services.
<code>LinkedAccount</code>	The linked account number. This is used for consolidated billing only. This dimension is included only for accounts that are linked to a separate paying account in a consolidated billing relationship. It is not included for accounts that are not linked to a consolidated billing paying account.
<code>Currency</code>	The monetary currency to bill the account. This dimension is required. Unit: USD

Amazon CloudFront Metrics and Dimensions

Amazon CloudFront sends metrics to Amazon CloudWatch for web distributions. Metrics and dimensions are not available for RTMP distributions. For more information, see [Monitoring CloudFront Activity Using CloudWatch](#) in the *Amazon CloudFront Developer Guide*.

Amazon CloudFront Metrics

The `AWS/CloudFront` namespace includes the following metrics.

Note

Only one statistic, Average or Sum, is applicable for each metric. However, all statistics are available through the console, API, and AWS Command Line Interface. In the following table, each metric specifies the statistic that is applicable to that metric.

Metric	Description
Requests	The number of requests for all HTTP methods and for both HTTP and HTTPS requests. Valid Statistics: Sum Units: Count
BytesDownloaded	The number of bytes downloaded by viewers for GET, HEAD, and OPTIONS requests. Valid Statistics: Sum Units: Bytes
BytesUploaded	The number of bytes uploaded to your origin with CloudFront using POST and PUT requests. Valid Statistics: Sum Units: Bytes
TotalErrorRate	The percentage of all requests for which the HTTP status code is 4xx or 5xx. Valid Statistics: Average Units: Percent
4xxErrorRate	The percentage of all requests for which the HTTP status code is 4xx. Valid Statistics: Average Units: Percent
5xxErrorRate	The percentage of all requests for which the HTTP status code is 5xx. Valid Statistics: Average Units: Percent

Dimensions for CloudFront Metrics

CloudFront metrics use the CloudFront namespace and provide metrics for two dimensions:

Dimension	Description
DistributionId	The CloudFront ID of the distribution for which you want to display metrics.
Region	The region for which you want to display metrics. This value must be Global. The Region dimension is

Dimension	Description
	different from the region in which CloudFront metrics are stored, which is US East (N. Virginia).

Amazon CloudSearch Metrics and Dimensions

Amazon CloudSearch sends metrics to Amazon CloudWatch. For more information, see [Monitoring an Amazon CloudSearch Domain with Amazon CloudWatch](#) in the *Amazon CloudSearch Developer Guide*.

Amazon CloudSearch Metrics

The AWS/CloudSearch namespace includes the following metrics.

Metric	Description
SuccessfulRequests	The number of search requests successfully processed by a search instance. Units: Count Valid statistics: Maximum, Sum
SearchableDocuments	The number of searchable documents in the domain's search index. Units: Count Valid statistics: Maximum
IndexUtilization	The percentage of the search instance's index capacity that has been used. The Maximum value indicates the percentage of the domain's index capacity that has been used. Units: Percent Valid statistics: Average, Maximum
Partitions	The number of partitions the index is distributed across. Units: Count Valid statistics: Minimum, Maximum

Dimensions for Amazon CloudSearch Metrics

Amazon CloudSearch sends the ClientId and DomainName dimensions to CloudWatch.

Dimension	Description
ClientId	The AWS account ID.
DomainName	The name of the search domain.

Amazon CloudWatch Events Metrics and Dimensions

CloudWatch Events sends metrics to Amazon CloudWatch every minute.

CloudWatch Events Metrics

The `AWS/Events` namespace includes the following metrics.

Metric	Description
<code>Invocations</code>	<p>Measures the number of times a target is invoked for a rule in response to an event. This includes successful and failed invocations, but does not include throttled or retried attempts until they fail permanently.</p> <p>Note CloudWatch Events only sends this metric to CloudWatch if it has a non-zero value.</p> <p>Valid Dimensions: RuleName</p> <p>Units: Count</p>
<code>FailedInvocations</code>	<p>Measures the number of invocations that failed permanently. This does not include invocations that are retried or that succeeded after a retry attempt.</p> <p>Valid Dimensions: RuleName</p> <p>Units: Count</p>
<code>TriggeredRules</code>	<p>Measures the number of triggered rules that matched with any event.</p> <p>Valid Dimensions: RuleName</p> <p>Units: Count</p>
<code>MatchedEvents</code>	<p>Measures the number of events that matched with any rule.</p> <p>Valid Dimensions: None</p> <p>Units: Count</p>
<code>ThrottledRules</code>	<p>Measures the number of triggered rules that are being throttled.</p> <p>Valid Dimensions: RuleName</p> <p>Units: Count</p>

Dimensions for CloudWatch Events Metrics

CloudWatch Events metrics have one dimension, which is listed below.

Dimension	Description
<code>RuleName</code>	Filters the available metrics by rule name.

Amazon CloudWatch Logs Metrics and Dimensions

CloudWatch Logs sends metrics to CloudWatch every minute.

CloudWatch Logs Metrics

The `AWS/Logs` namespace includes the following metrics.

Metric	Description
<code>IncomingBytes</code>	<p>The volume of log events in uncompressed bytes uploaded to CloudWatch Logs. When used with the <code>LogGroupName</code> dimension, this is the volume of log events in uncompressed bytes uploaded to the log group.</p> <p>Valid Dimensions: <code>LogGroupName</code></p> <p>Valid Statistic: Sum</p> <p>Units: Bytes</p>
<code>IncomingLogEvents</code>	<p>The number of log events uploaded to CloudWatch Logs. When used with the <code>LogGroupName</code> dimension, this is the number of log events uploaded to the log group.</p> <p>Valid Dimensions: <code>LogGroupName</code></p> <p>Valid Statistic: Sum</p> <p>Units: None</p>
<code>ForwardedBytes</code>	<p>The volume of log events in compressed bytes forwarded to the subscription destination.</p> <p>Valid Dimensions: <code>LogGroupName</code>, <code>DestinationType</code>, <code>FilterName</code></p> <p>Valid Statistic: Sum</p> <p>Units: Bytes</p>
<code>ForwardedLogEvents</code>	<p>The number of log events forwarded to the subscription destination.</p> <p>Valid Dimensions: <code>LogGroupName</code>, <code>DestinationType</code>, <code>FilterName</code></p> <p>Valid Statistic: Sum</p> <p>Units: None</p>
<code>DeliveryErrors</code>	<p>The number of log events for which CloudWatch Logs received an error when forwarding data to the subscription destination.</p> <p>Valid Dimensions: <code>LogGroupName</code>, <code>DestinationType</code>, <code>FilterName</code></p> <p>Valid Statistic: Sum</p> <p>Units: None</p>
<code>DeliveryThrottling</code>	<p>The number of log events for which CloudWatch Logs was throttled when forwarding data to the subscription destination.</p>

Metric	Description
	Valid Dimensions: LogGroupName, DestinationType, FilterName Valid Statistic: Sum Units: None

Dimensions for CloudWatch Logs Metrics

The dimensions that you can use with CloudWatch Logs metrics are listed below.

Dimension	Description
LogGroupName	The name of the CloudWatch Logs log group for which to display metrics.
DestinationType	The subscription destination for the CloudWatch Logs data, which can be AWS Lambda, Amazon Kinesis Streams, or Amazon Kinesis Firehose.
FilterName	The name of the subscription filter that is forwarding data from the log group to the destination. The subscription filter name is automatically converted by CloudWatch to ASCII and any unsupported characters get replaced with a question mark (?).

Amazon DynamoDB Metrics and Dimensions

Amazon DynamoDB sends metrics to CloudWatch. For more information, see [Monitoring DynamoDB Tables with Amazon CloudWatch](#) in the *Amazon DynamoDB Developer Guide*.

DynamoDB Metrics

The following metrics are available from Amazon DynamoDB. Note that DynamoDB only sends metrics to CloudWatch when they have a non-zero value. For example, the `UserErrors` metric is incremented whenever a request generates an HTTP 400 status code. If no HTTP 400 errors were encountered during a time period, CloudWatch will not provide metrics for `UserErrors` during that period.

Note

Amazon CloudWatch aggregates the following DynamoDB metrics at one-minute intervals:

- `ConditionalCheckFailedRequests`
- `ConsumedReadCapacityUnits`
- `ConsumedWriteCapacityUnits`
- `ReadThrottleEvents`
- `ReturnedBytes`
- `ReturnedItemCount`
- `ReturnedRecordsCount`
- `SuccessfulRequestLatency`
- `SystemErrors`
- `ThrottledRequests`
- `UserErrors`
- `WriteThrottleEvents`

For all other DynamoDB metrics, the aggregation granularity is five minutes.

Not all statistics, such as *Average* or *Sum*, are applicable for every metric. However, all of these values are available through the Amazon DynamoDB console, or by using the CloudWatch console, AWS CLI, or AWS SDKs for all metrics. In the following table, each metric has a list of Valid Statistics that is applicable to that metric.

Metric	Description
ConditionalCheckFailedRequests	<p>The number of failed attempts to perform conditional writes. The <code>PutItem</code>, <code>UpdateItem</code>, and <code>DeleteItem</code> operations let you provide a logical condition that must evaluate to true before the operation can proceed. If this condition evaluates to false, <code>ConditionalCheckFailedRequests</code> is incremented by one.</p> <p>Note A failed conditional write will result in an HTTP 400 error (Bad Request). These events are reflected in the <code>ConditionalCheckFailedRequests</code> metric, but not in the <code>UserErrors</code> metric.</p> <p>Units: Count</p> <p>Dimensions: <code>TableName</code></p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Minimum • Maximum • Average • <code>SampleCount</code> • Sum
ConsumedReadCapacityUnits	<p>The number of read capacity units consumed over the specified time period, so you can track how much of your provisioned throughput is used. You can retrieve the total consumed read capacity for a table and all of its global secondary indexes, or for a particular global secondary index. For more information, see Provisioned Throughput in Amazon DynamoDB.</p> <p>Note Use the <code>Sum</code> statistic to calculate the consumed throughput. For example, get the <code>Sum</code> value over a span of one minute, and divide it by the number of seconds in a minute (60) to calculate the average <code>ConsumedReadCapacityUnits</code> per second (recognizing that this average will not highlight any large but brief spikes in read activity that occurred during that minute). You can compare the calculated value to the provisioned throughput value you provide DynamoDB.</p> <p>Units: Count</p> <p>Dimensions: <code>TableName</code>, <code>GlobalSecondaryIndexName</code></p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • <code>Minimum</code> – Minimum number of read capacity units consumed by any individual request to the table or index.

Metric	Description
	<ul style="list-style-type: none"> • Maximum – Maximum number of read capacity units consumed by any individual request to the table or index. • Average – Average per-request read capacity consumed. • Sum – Total read capacity units consumed. This is the most useful statistic for the <code>ConsumedReadCapacityUnits</code> metric. • SampleCount – Number of requests to DynamoDB that consumed read capacity.
<code>ConsumedWriteCapacityUnits</code>	<p>The number of write capacity units consumed over the specified time period, so you can track how much of your provisioned throughput is used. You can retrieve the total consumed write capacity for a table and all of its global secondary indexes, or for a particular global secondary index. For more information, see Provisioned Throughput in Amazon DynamoDB.</p> <p>Note Use the <code>Sum</code> statistic to calculate the consumed throughput. For example, get the <code>Sum</code> value over a span of one minute, and divide it by the number of seconds in a minute (60) to calculate the average <code>ConsumedWriteCapacityUnits</code> per second (recognizing that this average will not highlight any large but brief spikes in write activity that occurred during that minute). You can compare the calculated value to the provisioned throughput value you provide DynamoDB.</p> <p>Units: Count</p> <p>Dimensions: <code>TableName</code>, <code>GlobalSecondaryIndexName</code></p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Minimum – Minimum number of write capacity units consumed by any individual request to the table or index. • Maximum – Maximum number of write capacity units consumed by any individual request to the table or index. • Average – Average per-request write capacity consumed. • Sum – Total write capacity units consumed. This is the most useful statistic for the <code>ConsumedWriteCapacityUnits</code> metric. • SampleCount – Number of requests to DynamoDB that consumed write capacity.

Metric	Description
OnlineIndexConsumedWriteCapacityUnits	<p>The number of write capacity units consumed when adding a new global secondary index to a table. If the write capacity of the index is too low, incoming write activity during the backfill phase might be throttled; this can increase the time it takes to create the index. You should monitor this statistic while the index is being built to determine whether the write capacity of the index is underprovisioned.</p> <p>You can adjust the write capacity of the index using the <code>UpdateTable</code> operation, even while the index is still being built.</p> <p>Note that the <code>ConsumedWriteCapacityUnits</code> metric for the index does not include the write throughput consumed during index creation.</p> <p>Units: Count</p> <p>Dimensions: <code>TableName</code>, <code>GlobalSecondaryIndexName</code></p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Minimum • Maximum • Average • SampleCount • Sum
OnlineIndexPercentageProgress	<p>The percentage of completion when a new global secondary index is being added to a table. DynamoDB must first allocate resources for the new index, and then backfill attributes from the table into the index. For large tables, this process might take a long time. You should monitor this statistic to view the relative progress as DynamoDB builds the index.</p> <p>Units: Count</p> <p>Dimensions: <code>TableName</code>, <code>GlobalSecondaryIndexName</code></p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Minimum • Maximum • Average • SampleCount • Sum

Metric	Description
OnlineIndexThrottleEvents	<p>The number of write throttle events that occur when adding a new global secondary index to a table. These events indicate that the index creation will take longer to complete, because incoming write activity is exceeding the provisioned write throughput of the index.</p> <p>You can adjust the write capacity of the index using the <code>UpdateTable</code> operation, even while the index is still being built.</p> <p>Note that the <code>WriteThrottleEvents</code> metric for the index does not include any throttle events that occur during index creation.</p> <p>Units: Count</p> <p>Dimensions: <code>TableName</code>, <code>GlobalSecondaryIndexName</code></p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Minimum • Maximum • Average • SampleCount • Sum
ProvisionedReadCapacityUnits	<p>The number of provisioned read capacity units for a table or a global secondary index.</p> <p>The <code>TableName</code> dimension returns the <code>ProvisionedReadCapacityUnits</code> for the table, but not for any global secondary indexes. To view <code>ProvisionedReadCapacityUnits</code> for a global secondary index, you must specify both <code>TableName</code> and <code>GlobalSecondaryIndex</code>.</p> <p>Units: Count</p> <p>Dimensions: <code>TableName</code>, <code>GlobalSecondaryIndexName</code></p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Minimum – Lowest setting for provisioned read capacity. If you use <code>UpdateTable</code> to increase read capacity, this metric shows the lowest value of provisioned <code>ReadCapacityUnits</code> during this time period. • Maximum – Highest setting for provisioned read capacity. If you use <code>UpdateTable</code> to decrease read capacity, this metric shows the highest value of provisioned <code>ReadCapacityUnits</code> during this time period. • Average – Average provisioned read capacity. The <code>ProvisionedReadCapacityUnits</code> metric is published at five-minute intervals. Therefore, if you rapidly adjust the provisioned read capacity units, this statistic might not reflect the true average.

Metric	Description
ProvisionedWriteCapacityUnits	<p>The number of provisioned write capacity units for a table or a global secondary index</p> <p>The <code>TableName</code> dimension returns the <code>ProvisionedWriteCapacityUnits</code> for the table, but not for any global secondary indexes. To view <code>ProvisionedWriteCapacityUnits</code> for a global secondary index, you must specify both <code>TableName</code> and <code>GlobalSecondaryIndex</code>.</p> <p>Units: Count</p> <p>Dimensions: <code>TableName</code>, <code>GlobalSecondaryIndexName</code></p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Minimum – Lowest setting for provisioned write capacity. If you use <code>UpdateTable</code> to increase write capacity, this metric shows the lowest value of provisioned <code>WriteCapacityUnits</code> during this time period. • Maximum – Highest setting for provisioned write capacity. If you use <code>UpdateTable</code> to decrease write capacity, this metric shows the highest value of provisioned <code>WriteCapacityUnits</code> during this time period. • Average – Average provisioned write capacity. The <code>ProvisionedWriteCapacityUnits</code> metric is published at five-minute intervals. Therefore, if you rapidly adjust the provisioned write capacity units, this statistic might not reflect the true average.
ReadThrottleEvents	<p>Requests to DynamoDB that exceed the provisioned read capacity units for a table or a global secondary index.</p> <p>A single request can result in multiple events. For example, a <code>BatchGetItem</code> that reads 10 items is processed as ten <code>GetItem</code> events. For each event, <code>ReadThrottleEvents</code> is incremented by one if that event is throttled. The <code>ThrottledRequests</code> metric for the entire <code>BatchGetItem</code> is not incremented unless <i>all ten</i> of the <code>GetItem</code> events are throttled.</p> <p>The <code>TableName</code> dimension returns the <code>ReadThrottleEvents</code> for the table, but not for any global secondary indexes. To view <code>ReadThrottleEvents</code> for a global secondary index, you must specify both <code>TableName</code> and <code>GlobalSecondaryIndex</code>.</p> <p>Units: Count</p> <p>Dimensions: <code>TableName</code>, <code>GlobalSecondaryIndexName</code></p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • <code>SampleCount</code> • <code>Sum</code>

Metric	Description
ReturnedBytes	<p>The number of bytes returned by <code>GetRecords</code> operations (Amazon DynamoDB Streams) during the specified time period.</p> <p>Units: Bytes</p> <p>Dimensions: Operation, StreamLabel, TableName</p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Minimum • Maximum • Average • SampleCount • Sum
ReturnedItemCount	<p>The number of items returned by <code>Query</code> or <code>Scan</code> operations during the specified time period.</p> <p>Note that the number of items <i>returned</i> is not necessarily the same as the number of items that were evaluated. For example, suppose you requested a <code>Scan</code> on a table that had 100 items, but specified a <code>FilterExpression</code> that narrowed the results so that only 15 items were returned. In this case, the response from <code>Scan</code> would contain a <code>ScanCount</code> of 100 and a <code>Count</code> of 15 returned items.</p> <p>Units: Count</p> <p>Dimensions: TableName</p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Minimum • Maximum • Average • SampleCount • Sum
ReturnedRecordsCount	<p>The number of stream records returned by <code>GetRecords</code> operations (Amazon DynamoDB Streams) during the specified time period.</p> <p>Units: Count</p> <p>Dimensions: Operation, StreamLabel, TableName</p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Minimum • Maximum • Average • SampleCount • Sum

Metric	Description
SuccessfulRequestLatency	<p>Successful requests to DynamoDB or Amazon DynamoDB Streams during the specified time period. SuccessfulRequestLatency can provide two different kinds of information:</p> <ul style="list-style-type: none"> • The elapsed time for successful requests (Minimum, Maximum, Sum, Or Average). • The number of successful requests (SampleCount). <p>SuccessfulRequestLatency reflects activity only within DynamoDB or Amazon DynamoDB Streams, and does not take into account network latency or client-side activity.</p> <p>Units: Milliseconds</p> <p>Dimensions: TableName, Operation</p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Minimum • Maximum • Average • SampleCount
SystemErrors	<p>Requests to DynamoDB or Amazon DynamoDB Streams that generate an HTTP 500 status code during the specified time period. An HTTP 500 usually indicates an internal service error.</p> <p>Units: Count</p> <p>Dimensions: All dimensions</p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Sum • SampleCount

Metric	Description
ThrottledRequests	<p>Requests to DynamoDB that exceed the provisioned throughput limits on a resource (such as a table or an index).</p> <p>ThrottledRequests is incremented by one if any event within a request exceeds a provisioned throughput limit. For example, if you update an item in a table with global secondary indexes, there are multiple events—a write to the table, and a write to each index. If one or more of these events are throttled, then ThrottledRequests is incremented by one.</p> <p>Note In a batch request (BatchGetItem or BatchWriteItem), ThrottledRequests is only incremented if every request in the batch is throttled. If any individual request within the batch is throttled, one of the following metrics is incremented:</p> <ul style="list-style-type: none"> • ReadThrottleEvents – For a throttled GetItem event within BatchGetItem. • WriteThrottleEvents – For a throttled PutItem or DeleteItem event within BatchWriteItem. <p>To gain insight into which event is throttling a request, compare ThrottledRequests with the ReadThrottleEvents and WriteThrottleEvents for the table and its indexes.</p> <p>Note A throttled request will result in an HTTP 400 status code. All such events are reflected in the ThrottledRequests metric, but not in the UserErrors metric.</p> <p>Units: Count</p> <p>Dimensions: TableName, Operation</p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Sum • SampleCount

Metric	Description
<p>UserErrors</p>	<p>Requests to DynamoDB or Amazon DynamoDB Streams that generate an HTTP 400 status code during the specified time period. An HTTP 400 usually indicates a client-side error such as an invalid combination of parameters, attempting to update a nonexistent table, or an incorrect request signature.</p> <p>All such events are reflected in the <code>UserErrors</code> metric, except for the following:</p> <ul style="list-style-type: none"> • <i>ProvisionedThroughputExceededException</i> – See the <code>ThrottledRequests</code> metric in this section. • <i>ConditionalCheckFailedException</i> – See the <code>ConditionalCheckFailedRequests</code> metric in this section. <p><code>UserErrors</code> represents the aggregate of HTTP 400 errors for DynamoDB or Amazon DynamoDB Streams requests for the current region and the current AWS account.</p> <p>Units: Count</p> <p>Dimensions: All dimensions</p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Sum • SampleCount
<p>WriteThrottleEvents</p>	<p>Requests to DynamoDB that exceed the provisioned write capacity units for a table or a global secondary index.</p> <p>A single request can result in multiple events. For example, a <code>PutItem</code> request on a table with three global secondary indexes would result in four events—the table write, and each of the three index writes. For each event, the <code>WriteThrottleEvents</code> metric is incremented by one if that event is throttled. For single <code>PutItem</code> requests, if any of the events are throttled, <code>ThrottledRequests</code> is also incremented by one. For <code>BatchWriteItem</code>, the <code>ThrottledRequests</code> metric for the entire <code>BatchWriteItem</code> is not incremented unless all of the individual <code>PutItem</code> or <code>DeleteItem</code> events are throttled.</p> <p>The <code>TableName</code> dimension returns the <code>WriteThrottleEvents</code> for the table, but not for any global secondary indexes. To view <code>WriteThrottleEvents</code> for a global secondary index, you must specify both <code>TableName</code> and <code>GlobalSecondaryIndexName</code>.</p> <p>Units: Count</p> <p>Dimensions: <code>TableName</code>, <code>GlobalSecondaryIndexName</code></p> <p>Valid Statistics:</p> <ul style="list-style-type: none"> • Sum • SampleCount

Dimensions for DynamoDB Metrics

The metrics for DynamoDB are qualified by the values for the account, table name, global secondary index name, or operation. You can use the CloudWatch console to retrieve DynamoDB data along any of the dimensions in the table below.

Dimension	Description
GlobalSecondaryIndexName	This dimension limits the data to a global secondary index on a table. If you specify <code>GlobalSecondaryIndexName</code> , you must also specify <code>TableName</code> .
Operation	<p>This dimension limits the data to one of the following DynamoDB operations:</p> <ul style="list-style-type: none"> • <code>PutItem</code> • <code>DeleteItem</code> • <code>UpdateItem</code> • <code>GetItem</code> • <code>BatchGetItem</code> • <code>Scan</code> • <code>Query</code> • <code>BatchWriteItem</code> <p>In addition, you can limit the data to the following Amazon DynamoDB Streams operation:</p> <ul style="list-style-type: none"> • <code>GetRecords</code>
StreamLabel	This dimension limits the data to a specific stream label. It is used with metrics originating from Amazon DynamoDB Streams <code>GetRecords</code> operations.
TableName	This dimension limits the data to a specific table. This value can be any table name in the current region and the current AWS account.

Amazon EC2 Metrics and Dimensions

Amazon Elastic Compute Cloud (Amazon EC2) sends metrics to CloudWatch for your EC2 instances. Basic (five-minute) monitoring is enabled by default. You can enable detailed (one-minute) monitoring. For information about additional metrics for Amazon EC2 instances that are in an Auto Scaling group, see [Auto Scaling Metrics and Dimensions \(p. 50\)](#).

For more information about how to monitor Amazon EC2, see [Monitoring Your Instances with CloudWatch](#) in the *Amazon EC2 User Guide for Linux Instances*.

Amazon EC2 Metrics

The following metrics are available from each EC2 instance.

The `AWS/EC2` namespace includes the following CPU credit metrics for your T2 instances. CPU credit metrics are available at a 5 minute frequency.

Metric	Description
CPUCreditUsage	<p>[T2 instances] The number of CPU credits consumed during the specified period.</p> <p>This metric identifies the amount of time during which physical CPUs were used for processing instructions by virtual CPUs allocated to the instance.</p> <p>CPU Credit metrics are available at a 5 minute frequency.</p> <p>Units: Count</p>
CPUCreditBalance	<p>[T2 instances] The number of CPU credits that an instance has accumulated.</p> <p>This metric determines how long an instance can burst beyond its baseline performance level at a given rate.</p> <p>CPU Credit metrics are available at a 5 minute frequency.</p> <p>Units: Count</p>

The `AWS/EC2` namespace includes the following instance metrics.

Metric	Description
CPUUtilization	<p>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 upon a selected instance.</p> <p>To use the percentiles statistic, you must enable detailed monitoring.</p> <p>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.</p> <p>Units: Percent</p>
DiskReadOps	<p>Completed read operations from all instance store volumes available to the instance in a specified period of time.</p> <p>To calculate the average I/O operations per second (IOPS) for the period, divide the total operations in the period by the number of seconds in that period.</p> <p>Units: Count</p>
DiskWriteOps	<p>Completed write operations to all instance store volumes available to the instance in a specified period of time.</p> <p>To calculate the average I/O operations per second (IOPS) for the period, divide the total operations in the period by the number of seconds in that period.</p> <p>Units: Count</p>
DiskReadBytes	<p>Bytes read from all instance store volumes available to the instance.</p>

Metric	Description
	<p>This metric is used to determine the volume of the data the application reads from the hard disk of the instance. This can be used to determine the speed of the application.</p> <p>Units: Bytes</p>
DiskWriteBytes	<p>Bytes written to all instance store volumes available to the instance.</p> <p>This metric is used to determine the volume of the data the application writes onto the hard disk of the instance. This can be used to determine the speed of the application.</p> <p>Units: Bytes</p>
NetworkIn	<p>The number of bytes received on all network interfaces by the instance. This metric identifies the volume of incoming network traffic to an application on a single instance.</p> <p>Units: Bytes</p>
NetworkOut	<p>The number of bytes sent out on all network interfaces by the instance. This metric identifies the volume of outgoing network traffic to an application on a single instance.</p> <p>Units: Bytes</p>
NetworkPacketsIn	<p>The number of packets received on all network interfaces by the instance. This metric identifies the volume of incoming traffic in terms of the number of packets on a single instance. This metric is available for basic monitoring only.</p> <p>Units: Count</p> <p>Statistics: Minimum, Maximum, Average</p>
NetworkPacketsOut	<p>The number of packets sent out on all network interfaces by the instance. This metric identifies the volume of outgoing traffic in terms of the number of packets on a single instance. This metric is available for basic monitoring only.</p> <p>Units: Count</p> <p>Statistics: Minimum, Maximum, Average</p>

The `AWS/EC2` namespace includes the following status checks metrics. Status check metrics are available at a 1 minute frequency. For a newly-launched instance, status check metric data is only available after the instance has completed the initialization state (within a few minutes of the instance entering the running state).

Metric	Description
StatusCheckFailed	<p>Reports whether the instance has passed both the instance status check and the system status check in the last minute.</p> <p>This metric can be either 0 (passed) or 1 (failed).</p> <p>Units: Count</p>

Metric	Description
StatusCheckFailed_Instance	<p>Reports whether the instance has passed the instance status check in the last minute.</p> <p>This metric can be either 0 (passed) or 1 (failed).</p> <p>Units: Count</p>
StatusCheckFailed_System	<p>Reports whether the instance has passed the system status check in the last minute.</p> <p>This metric can be either 0 (passed) or 1 (failed).</p> <p>Units: Count</p>

Amazon CloudWatch data for a new EC2 instance typically becomes available within one minute of the end of the first period of time requested (the *aggregation period*) in the query. You can set the period—the length of time over which statistics are aggregated—with the `Period` parameter. For more information on periods, see [Periods \(p. 6\)](#).

You can use the currently available dimensions for EC2 instances (for example, `ImageId` or `InstanceType`) to refine the metrics returned. For information about the dimensions you can use with EC2, see [Dimensions for Amazon EC2 Metrics \(p. 69\)](#).

Dimensions for Amazon EC2 Metrics

If you're using Detailed Monitoring, you can filter the EC2 instance data using any of the dimensions in the following table.

Dimension	Description
AutoScalingGroupName	This dimension filters the data you request for all instances in a specified capacity group. An <i>Auto Scaling group</i> is a collection of instances you define if you're using Auto Scaling. This dimension is available only for Amazon EC2 metrics when the instances are in such an Auto Scaling group. Available for instances with Detailed or Basic Monitoring enabled.
ImageId	This dimension filters the data you request for all instances running this Amazon EC2 Amazon Machine Image (AMI). Available for instances with Detailed Monitoring enabled.
InstanceId	This dimension filters the data you request for the identified instance only. This helps you pinpoint an exact instance from which to monitor data.
InstanceType	This dimension filters the data you request for all instances running with this specified instance type. This helps you categorize your data by the type of instance running. For example, you might compare data from an <code>m1.small</code> instance and an <code>m1.large</code> instance to determine which has the better business value for your application. Available for instances with Detailed Monitoring enabled.

Amazon EC2 Spot Fleet Metrics and Dimensions

Amazon Elastic Compute Cloud (Amazon EC2) sends information about your Spot fleet to CloudWatch. For more information, see [CloudWatch Metrics for Spot Fleet](#) in the *Amazon EC2 User Guide for Linux Instances*.

Amazon EC2 Spot Fleet Metrics

The `AWS/EC2Spot` namespace includes the following metrics, plus the CloudWatch metrics for the Spot instances in your fleet.

The `AWS/EC2Spot` namespace includes the following metrics.

Metric	Description
<code>AvailableInstancePoolsCount</code>	The Spot Instance pools specified in the Spot Fleet request. Units: Count
<code>BidsSubmittedForCapacity</code>	The capacity for which Amazon EC2 has submitted bids. Units: Count
<code>EligibleInstancePoolCount</code>	The Spot Instance pools specified in the Spot Fleet request where Amazon EC2 can fulfill bids. Amazon EC2 will not fulfill bids in pools where your bid price is less than the Spot price or the Spot price is greater than the price for On-Demand instances. Units: Count
<code>FulfilledCapacity</code>	The capacity that Amazon EC2 has fulfilled. Units: Count
<code>MaxPercentCapacityAllocation</code>	The maximum value of <code>PercentCapacityAllocation</code> across all Spot Instance pools specified in the Spot Fleet request. Units: Percent
<code>PendingCapacity</code>	The difference between <code>TargetCapacity</code> and <code>FulfilledCapacity</code> . Units: Count
<code>PercentCapacityAllocation</code>	The capacity allocated for the Spot Instance pool for the specified dimensions. To get the maximum value recorded across all Spot Instance pools, use <code>MaxPercentCapacityAllocation</code> . Units: Percent
<code>TargetCapacity</code>	The target capacity of the Spot Fleet request. Units: Count
<code>TerminatingCapacity</code>	The capacity that is being terminated due to Spot Instance interruptions. Units: Count

If the unit of measure for a metric is `Count`, the most useful statistic is `Average`.

Dimensions for Amazon EC2 Spot Fleet Metrics

You can filter the data using the following dimensions.

Dimensions	Description
<code>AvailabilityZone</code>	Filter the data by Availability Zone.
<code>FleetRequestId</code>	Filter the data by Spot Fleet request.
<code>InstanceType</code>	Filter the data by instance type.

Amazon ECS Metrics and Dimensions

Amazon EC2 Container Service (Amazon ECS) sends metrics to Amazon CloudWatch. For more information, see [Amazon ECS CloudWatch Metrics](#) in the *Amazon EC2 Container Service Developer Guide*.

Amazon ECS Metrics

Amazon ECS provides metrics for you to monitor the CPU and memory reservation and utilization across your cluster as a whole, and the CPU and memory utilization on the services in your clusters.

Amazon ECS sends the following metrics to CloudWatch every minute.

Metric	Description
<code>CPUReservation</code>	<p>The percentage of CPU units that are reserved by running tasks in the cluster.</p> <p>Cluster CPU reservation (this metric can only be filtered by <code>ClusterName</code>) is measured as the total CPU units that are reserved by Amazon ECS tasks on the cluster, divided by the total CPU units that were registered for all of the container instances in the cluster.</p> <p>Valid Dimensions: <code>ClusterName</code>, <code>ServiceName</code></p> <p>Valid Statistics: <code>Average</code>, <code>Minimum</code>, <code>Maximum</code>, <code>Sum</code>, <code>Data Samples</code>.</p> <p>Unit: <code>Percent</code></p>
<code>CPUUtilization</code>	<p>The percentage of CPU units that are used in the cluster or service.</p> <p>Cluster CPU utilization (metrics that are filtered by <code>ClusterName</code> without <code>ServiceName</code>) is measured as the total CPU units in use by Amazon ECS tasks on the cluster, divided by the total CPU units that were registered for all of the container instances in the cluster.</p> <p>Service CPU utilization (metrics that are filtered by <code>ClusterName</code> and <code>ServiceName</code>) is measured as the</p>

Metric	Description
	<p>total CPU units in use by the tasks that belong to the service, divided by the total number of CPU units that are reserved for the tasks that belong to the service.</p> <p>Valid Dimensions: <code>ClusterName</code>, <code>ServiceName</code></p> <p>Valid Statistics: Average, Minimum, Maximum, Sum, Data Samples.</p> <p>Unit: Percent</p>
MemoryReservation	<p>The percentage of memory that is reserved by running tasks in the cluster.</p> <p>Cluster memory reservation (this metric can only be filtered by <code>ClusterName</code>) is measured as the total memory that is reserved by Amazon ECS tasks on the cluster, divided by the total amount of memory that was registered for all of the container instances in the cluster.</p> <p>Valid Dimensions: <code>ClusterName</code>, <code>ServiceName</code></p> <p>Valid Statistics: Average, Minimum, Maximum, Sum, Data Samples.</p> <p>Unit: Percent</p>
MemoryUtilization	<p>The percentage of memory that is used in the cluster or service.</p> <p>Cluster memory utilization (metrics that are filtered by <code>ClusterName</code> without <code>ServiceName</code>) is measured as the total memory in use by Amazon ECS tasks on the cluster, divided by the total amount of memory that was registered for all of the container instances in the cluster.</p> <p>Service memory utilization (metrics that are filtered by <code>ClusterName</code> and <code>ServiceName</code>) is measured as the total memory in use by the tasks that belong to the service, divided by the total memory that is reserved for the tasks that belong to the service.</p> <p>Valid Dimensions: <code>ClusterName</code>, <code>ServiceName</code></p> <p>Valid Statistics: Average, Minimum, Maximum, Sum, Data Samples.</p> <p>Unit: Percent</p>

Dimensions for Amazon ECS Metrics

Amazon ECS metrics use the `AWS/ECS` namespace and provide metrics for the following dimensions:

Dimension	Description
ClusterName	This dimension filters the data you request for all resources in a specified cluster. All Amazon ECS metrics are filtered by ClusterName.
ServiceName	This dimension filters the data you request for all resources in a specified service within a specified cluster.

AWS Elastic Beanstalk Metrics and Dimensions

AWS Elastic Beanstalk sends metrics to Amazon CloudWatch. For more information, see [Publishing Amazon CloudWatch Custom Metrics for an Environment](#) in the *AWS Elastic Beanstalk Developer Guide*.

Elastic Beanstalk Metrics

The AWS/ElasticBeanstalk namespace includes the following metrics.

Metric	Description
EnvironmentHealth	[Environment] The health status of the environment. The possible values are 0 (OK), 1 (Info), 5 (Unknown), 10 (No data), 15 (Warning), 20 (Degraded) and 25 (Severe).
InstancesOk	[Environment] The number of instances with OK health status.
InstancesPending	[Environment] The number of instances with Pending health status.
InstancesInfo	[Environment] The number of instances with Info health status.
InstancesUnknown	[Environment] The number of instances with Unknown health status.
InstancesNoData	[Environment] The number of instances with no health status data.
InstancesWarning	[Environment] The number of instances with Warning health status.
InstancesDegraded	[Environment] The number of instances with Degraded health status.
InstancesSevere	[Environment] The number of instances with Severe health status.
ApplicationRequestsTotal	The number of requests completed by the instance or environment.
ApplicationRequests2xx	The number of requests that completed with a 2XX status code.
ApplicationRequests3xx	The number of requests that completed with a 3XX status code.
ApplicationRequests4xx	The number of requests that completed with a 4XX status code.
ApplicationRequests5xx	The number of requests that completed with a 5XX status code.
ApplicationLatencyP10	The average time to complete the fastest 10 percent of requests.

Metric	Description
ApplicationLatencyP50	The average time to complete the fastest 50 percent of requests.
ApplicationLatencyP75	The average time to complete the fastest 75 percent of requests.
ApplicationLatencyP85	The average time to complete the fastest 85 percent of requests.
ApplicationLatencyP90	The average time to complete the fastest 90 percent of requests.
ApplicationLatencyP95	The average time to complete the fastest 95 percent of requests.
ApplicationLatencyP99	The average time to complete the fastest 99 percent of requests.
ApplicationLatencyP99.9	The average time to complete the fastest x percent of requests.
LoadAverage1min	[Instance] The average CPU load over the last minute.
InstanceHealth	[Instance] The health status of the instance.
RootFilesystemUtil	[Instance] The percentage of disk space in use.
CPUIrq	[Instance] The percentage of time the CPU was in this state in the last minute.
CPUser	[Instance] The percentage of time the CPU was in this state in the last minute.
CPUIidle	[Instance] The percentage of time the CPU was in this state in the last minute.
CPUSystem	[Instance] The percentage of time the CPU was in this state in the last minute.
CPUSoftirq	[Instance] The percentage of time the CPU was in this state in the last minute.
CPUIowait	[Instance] The percentage of time the CPU was in this state in the last minute.
CPUNice	[Instance] The percentage of time the CPU was in this state in the last minute.

Dimensions for Elastic Beanstalk Metrics

You can filter the data using the following dimensions.

Dimensions	Description
Environment	Filter the data by environment.
Instance	Filter the data by instance.

Amazon ElastiCache Metrics and Dimensions

Amazon ElastiCache sends metrics to Amazon CloudWatch. For more information, see [Viewing Cache Cluster and Cache Node Metrics](#) in the *Amazon ElastiCache User Guide*.

Contents

- [Dimensions for ElastiCache Metrics \(p. 75\)](#)
- [Host-Level Metrics \(p. 75\)](#)
- [Metrics for Memcached \(p. 76\)](#)
- [Metrics for Redis \(p. 78\)](#)

Dimensions for ElastiCache Metrics

All ElastiCache metrics use the `AWS/ElastiCache` namespace and provide metrics for a single dimension, the `CacheNodeId`, which is the automatically-generated identifier for each cache node in the cache cluster. You can find out what these values are for your cache nodes by using the `DescribeCacheClusters` API or **describe-cache-clusters** command line utility. For more information, see [DescribeCacheClusters](#) in the *Amazon ElastiCache API Reference* and [describe-cache-clusters](#) in the *AWS Command Line Interface Reference*.

Each metric is published under a single set of dimensions. When retrieving metrics, you must supply both the `CacheClusterId` and `CacheNodeId` dimensions.

Topics

- [Host-Level Metrics \(p. 75\)](#)
- [Metrics for Memcached \(p. 76\)](#)
- [Metrics for Redis \(p. 78\)](#)
- [Which Metrics Should I Monitor?](#)

Host-Level Metrics

The `AWS/ElastiCache` namespace includes the following host-level metrics for individual cache nodes.

See Also

- [Metrics for Memcached \(p. 76\)](#)
- [Metrics for Redis \(p. 78\)](#)

Metric	Description	Unit
<code>CPUtilization</code>	The percentage of CPU utilization.	Percent
<code>FreeableMemory</code>	The amount of free memory available on the host.	Bytes
<code>NetworkBytesIn</code>	The number of bytes the host has read from the network.	Bytes
<code>NetworkBytesOut</code>	The number of bytes the host has written to the network.	Bytes
<code>SwapUsage</code>	The amount of swap used on the host.	Bytes

Metrics for Memcached

The `AWS/ElastiCache` namespace includes the following metrics that are derived from the Memcached `stats` command. Each metric is calculated at the cache node level.

For complete documentation of the Memcached `stats` command, go to <https://github.com/memcached/memcached/blob/master/doc/protocol.txt>.

See Also

- [Host-Level Metrics \(p. 75\)](#)

Metric	Description	Unit
BytesReadIntoMemcached	The number of bytes that have been read from the network by the cache node.	Bytes
BytesUsedForCacheItems	The number of bytes used to store cache items.	Bytes
BytesWrittenOutFromMemcached	The number of bytes that have been written to the network by the cache node.	Bytes
CasBadval	The number of CAS (check and set) requests the cache has received where the Cas value did not match the Cas value stored.	Count
CasHits	The number of Cas requests the cache has received where the requested key was found and the Cas value matched.	Count
CasMisses	The number of Cas requests the cache has received where the key requested was not found.	Count
CmdFlush	The number of flush commands the cache has received.	Count
CmdGet	The number of get commands the cache has received.	Count
CmdSet	The number of set commands the cache has received.	Count
CurrConnections	A count of the number of connections connected to the cache at an instant in time.	Count
CurrItems	A count of the number of items currently stored in the cache.	Count
DecrHits	The number of decrement requests the cache has received where the requested key was found.	Count
DecrMisses	The number of decrement requests the cache has received where the requested key was not found.	Count
DeleteHits	The number of delete requests the cache has received where the requested key was found.	Count

Metric	Description	Unit
DeleteMisses	The number of delete requests the cache has received where the requested key was not found.	Count
Evictions	The number of non-expired items the cache evicted to allow space for new writes.	Count
GetHits	The number of get requests the cache has received where the key requested was found.	Count
GetMisses	The number of get requests the cache has received where the key requested was not found.	Count
IncrHits	The number of increment requests the cache has received where the key requested was found.	Count
IncrMisses	The number of increment requests the cache has received where the key requested was not found.	Count
Reclaimed	The number of expired items the cache evicted to allow space for new writes.	Count

For Memcached 1.4.14, the following additional metrics are provided.

Metric	Description	Unit
BytesUsedForHash	The number of bytes currently used by hash tables.	Bytes
CmdConfigGet	The cumulative number of config get requests.	Count
CmdConfigSet	The cumulative number of config set requests.	Count
CmdTouch	The cumulative number of touch requests.	Count
CurrConfig	The current number of configurations stored.	Count
EvictedUnfetched	The number of valid items evicted from the least recently used cache (LRU) which were never touched after being set.	Count
ExpiredUnfetched	The number of expired items reclaimed from the LRU which were never touched after being set.	Count
SlabsMoved	The total number of slab pages that have been moved.	Count
TouchHits	The number of keys that have been touched and were given a new expiration time.	Count
TouchMisses	The number of items that have been touched, but were not found.	Count

The AWS/ElastiCache namespace includes the following calculated cache-level metrics.

Metric	Description	Unit
NewConnections	The number of new connections the cache has received. This is derived from the memcached	Count

Metric	Description	Unit
	<code>total_connections</code> statistic by recording the change in <code>total_connections</code> across a period of time. This will always be at least 1, due to a connection reserved for a ElastiCache.	
NewItems	The number of new items the cache has stored. This is derived from the memcached <code>total_items</code> statistic by recording the change in <code>total_items</code> across a period of time.	Count
UnusedMemory	<p>The amount of memory not used by data. This is derived from the Memcached statistics <code>limit_maxbytes</code> and <code>bytes</code> by subtracting <code>bytes</code> from <code>limit_maxbytes</code>.</p> <p>Because Memcached overhead uses memory in addition to that used by data, <code>UnusedMemory</code> should not be considered to be the amount of memory available for additional data. You may experience evictions even though you still have some unused memory.</p> <p>For more detailed information, see Memcached item memory usage.</p>	Bytes

Metrics for Redis

The `AWS/ElastiCache` namespace includes the following Redis metrics.

With the exception of `ReplicationLag`, these metrics are derived from the Redis `info` command. Each metric is calculated at the cache node level.

For complete documentation of the Redis `info` command, go to <http://redis.io/commands/info>.

See Also

- [Host-Level Metrics \(p. 75\)](#)

Metric	Description	Unit
BytesUsedForCache	The total number of bytes allocated by Redis.	Bytes
CacheHits	The number of successful key lookups.	Count
CacheMisses	The number of unsuccessful key lookups.	Count
CurrConnections	The number of client connections, excluding connections from read replicas.	Count
Evictions	The number of keys that have been evicted due to the <code>maxmemory</code> limit.	Count
HyperLogLogBasedCmds	The total number of HyperLogLog based commands. This is derived from the Redis <code>commandstats</code> statistic by summing all of the <code>pf</code> type of commands (<code>pfadd</code> , <code>pfcount</code> , <code>pfmerge</code>).	Count

Metric	Description	Unit
NewConnections	The total number of connections that have been accepted by the server during this period.	Count
Reclaimed	The total number of key expiration events.	Count
ReplicationBytes	For primaries with attached replicas, <code>ReplicationBytes</code> reports the number of bytes that the primary is sending to all of its replicas. This metric is representative of the write load on the replication group. For replicas and standalone primaries, <code>ReplicationBytes</code> is always 0.	Bytes
ReplicationLag	This metric is only applicable for a cache node running as a read replica. It represents how far behind, in seconds, the replica is in applying changes from the primary cache cluster.	Seconds
SaveInProgress	This binary metric returns 1 whenever a background save (forked or forkless) is in progress, and 0 otherwise. A background save process is typically used during snapshots and syncs. These operations can cause degraded performance. Using the <code>SaveInProgress</code> metric, you can diagnose whether or not degraded performance was caused by a background save process.	Count

These are aggregations of certain kinds of commands, derived from **info commandstats**:

Metric	Description	Unit
CurrItems	The number of items in the cache. This is derived from the Redis <code>keyspace</code> statistic, summing all of the keys in the entire keyspace.	Count
GetTypeCmds	The total number of get types of commands. This is derived from the Redis <code>commandstats</code> statistic by summing all of the get types of commands (get , mget , hget , etc.)	Count
HashBasedCmds	The total number of commands that are hash-based. This is derived from the Redis <code>commandstats</code> statistic by summing all of the commands that act upon one or more hashes.	Count
KeyBasedCmds	The total number of commands that are key-based. This is derived from the Redis <code>commandstats</code> statistic by summing all of the commands that act upon one or more keys.	Count
ListBasedCmds	The total number of commands that are list-based. This is derived from the Redis <code>commandstats</code> statistic by summing all of the commands that act upon one or more lists.	Count

Metric	Description	Unit
SetBasedCmds	The total number of commands that are set-based. This is derived from the Redis <code>commandstats</code> statistic by summing all of the commands that act upon one or more sets.	Count
SetTypeCmds	The total number of set types of commands. This is derived from the Redis <code>commandstats</code> statistic by summing all of the set types of commands (set , hset , etc.)	Count
SortedSetBasedCmds	The total number of commands that are sorted set-based. This is derived from the Redis <code>commandstats</code> statistic by summing all of the commands that act upon one or more sorted sets.	Count
StringBasedCmds	The total number of commands that are string-based. This is derived from the Redis <code>commandstats</code> statistic by summing all of the commands that act upon one or more strings.	Count

Amazon EBS Metrics and Dimensions

Amazon Elastic Block Store (Amazon EBS) sends data points to CloudWatch for several metrics. Amazon EBS General Purpose SSD (gp2), Throughput Optimized HDD (st1), Cold HDD (sc1), and Magnetic (standard) volumes automatically send five-minute metrics to CloudWatch. Provisioned IOPS SSD (io1) volumes automatically send one-minute metrics to CloudWatch. For more information, see [Monitoring the Status of Your Volumes](#) in the *Amazon EC2 User Guide for Linux Instances*.

Amazon EBS Metrics

Amazon Elastic Block Store (Amazon EBS) sends data points to CloudWatch for several metrics. Amazon EBS General Purpose SSD (gp2), Throughput Optimized HDD (st1), Cold HDD (sc1), and Magnetic (standard) volumes automatically send five-minute metrics to CloudWatch. Provisioned IOPS SSD (io1) volumes automatically send one-minute metrics to CloudWatch. For more information about how to monitor Amazon EBS, see [Monitoring the Status of Your Volumes](#) in the *Amazon EC2 User Guide for Linux Instances*.

The `AWS/EBS` namespace includes the following metrics.

Metric	Description
VolumeReadBytes VolumeWriteBytes	Provides information on the I/O operations in a specified period of time. The <code>Sum</code> statistic reports the total number of bytes transferred during the period. The <code>Average</code> statistic reports the average size of each I/O operation during the period. The <code>SampleCount</code> statistic reports the total number of I/O operations during the period. The <code>Minimum</code> and <code>Maximum</code> statistics are not relevant for this metric. Data is only reported to Amazon CloudWatch when the volume is active. If the volume is idle, no data is reported to Amazon CloudWatch. Units: Bytes
VolumeReadOps	The total number of I/O operations in a specified period of time.

Metric	Description
VolumeWriteOps	To calculate the average I/O operations per second (IOPS) for the period, divide the total operations in the period by the number of seconds in that period. Units: Count
VolumeTotalReadTime VolumeTotalWriteTime	The total number of seconds spent by all operations that completed in a specified period of time. If multiple requests are submitted at the same time, this total could be greater than the length of the period. For example, for a period of 5 minutes (300 seconds): if 700 operations completed during that period, and each operation took 1 second, the value would be 700 seconds. Units: Seconds
VolumeIdleTime	The total number of seconds in a specified period of time when no read or write operations were submitted. Units: Seconds
VolumeQueueLength	The number of read and write operation requests waiting to be completed in a specified period of time. Units: Count
VolumeThroughputPercentage	Used with Provisioned IOPS SSD volumes only. The percentage of I/O operations per second (IOPS) delivered of the total IOPS provisioned for an Amazon EBS volume. Provisioned IOPS SSD volumes deliver within 10 percent of the provisioned IOPS performance 99.9 percent of the time over a given year. During a write, if there are no other pending I/O requests in a minute, the metric value will be 100 percent. Also, a volume's I/O performance may become degraded temporarily due to an action you have taken (for example, creating a snapshot of a volume during peak usage, running the volume on a non-EBS-optimized instance, accessing data on the volume for the first time). Units: Percent
VolumeConsumedReadWriteOps	Used with Provisioned IOPS SSD volumes only. The total amount of read and write operations (normalized to 256K capacity units) consumed in a specified period of time. I/O operations that are smaller than 256K each count as 1 consumed IOPS. I/O operations that are larger than 256K are counted in 256K capacity units. For example, a 1024K I/O would count as 4 consumed IOPS. Units: Count
BurstBalance	Used with General Purpose SSD (gp2), Throughput Optimized HDD (st1), and Cold HDD (sc1) volumes only. Provides information about the percentage of I/O credits (for gp2) or throughput credits (for st1 and sc1) remaining in the burst bucket. Data is reported to CloudWatch only when the volume is active. If the volume is not attached, no data is reported. Units: Percent

Dimensions for Amazon EBS Metrics

The only dimension that Amazon EBS sends to CloudWatch is the Volume ID. This means that all available statistics are filtered by Volume ID.

Amazon EFS Metrics and Dimensions

Amazon EFS sends metrics to CloudWatch for every Amazon EFS file system every minute. For more information, see [Monitor Metrics with CloudWatch](#) in the *Amazon Elastic File System User Guide*.

Amazon CloudWatch Metrics for Amazon EFS

The `AWS/EFS` namespace includes the following metrics.

Metric	Description
<code>BurstCreditBalance</code>	<p>The number of burst credits that a file system has.</p> <p>Burst credits allow a file system to burst to throughput levels above a file system's baseline level for periods of time. For more information, see Throughput scaling in Amazon EFS.</p> <p>The <code>Minimum</code> statistic is the smallest burst credit balance for any minute during the period. The <code>Maximum</code> statistic is the largest burst credit balance for any minute during the period. The <code>Average</code> statistic is the average burst credit balance during the period.</p> <p>Units: Bytes</p> <p>Valid statistics: <code>Minimum</code>, <code>Maximum</code>, <code>Average</code></p>
<code>ClientConnections</code>	<p>The number of client connections to a file system. When using a standard client, there is one connection per mounted Amazon EC2 instance.</p> <p>Note To calculate the average <code>ClientConnections</code> for periods greater than one minute, divide the <code>Sum</code> statistic by the number of minutes in the period.</p> <p>Units: Count of client connections</p> <p>Valid statistics: <code>Sum</code></p>
<code>DataReadIOBytes</code>	<p>The number of bytes for each file system read operation.</p> <p>The <code>Sum</code> statistic is the total number of bytes associated with read operations. The <code>Minimum</code> statistic is the size of the smallest read operation during the period. The <code>Maximum</code> statistic is the size of the largest read operation during the period. The <code>Average</code> statistic is the average size of read operations during the period. The <code>SampleCount</code> statistic provides a count of read operations.</p> <p>Units:</p> <ul style="list-style-type: none">• Bytes for <code>Minimum</code>, <code>Maximum</code>, <code>Average</code>, and <code>Sum</code>.• Count for <code>SampleCount</code>.

Metric	Description
	Valid statistics: Minimum, Maximum, Average, Sum, SampleCount
DataWriteIOBytes	<p>The number of bytes for each file write operation.</p> <p>The <code>Sum</code> statistic is the total number of bytes associated with write operations. The <code>Minimum</code> statistic is the size of the smallest write operation during the period. The <code>Maximum</code> statistic is the size of the largest write operation during the period. The <code>Average</code> statistic is the average size of write operations during the period. The <code>SampleCount</code> statistic provides a count of write operations.</p> <p>Units:</p> <ul style="list-style-type: none"> • Bytes are the units for the <code>Minimum</code>, <code>Maximum</code>, <code>Average</code>, and <code>Sum</code> statistics. • Count for <code>SampleCount</code>. <p>Valid statistics: Minimum, Maximum, Average, Sum, SampleCount</p>
MetadataIOBytes	<p>The number of bytes for each metadata operation.</p> <p>The <code>Sum</code> statistic is the total number of bytes associated with metadata operations. The <code>Minimum</code> statistic is the size of the smallest metadata operation during the period. The <code>Maximum</code> statistic is the size of the largest metadata operation during the period. The <code>Average</code> statistic is the size of the average metadata operation during the period. The <code>SampleCount</code> statistic provides a count of metadata operations.</p> <p>Units:</p> <ul style="list-style-type: none"> • Bytes are the units for the <code>Minimum</code>, <code>Maximum</code>, <code>Average</code>, and <code>Sum</code> statistics. • Count for <code>SampleCount</code>. <p>Valid statistics: Minimum, Maximum, Average, Sum, SampleCount</p>
PercentIOLimit	<p>Shows how close a file system is to reaching the I/O limit of the General Purpose performance mode. If this metric is at 100% more often than not, consider moving your application to a file system using the Max I/O performance mode.</p> <p>Note This metric is only submitted for file systems using the General Purpose performance mode.</p> <p>Units:</p> <ul style="list-style-type: none"> • Percent

Metric	Description
PermittedThroughput	<p>The maximum amount of throughput a file system is allowed, given the file system size and <code>BurstCreditBalance</code>. For more information, see Amazon EFS Performance.</p> <p>The <code>Minimum</code> statistic is the smallest throughput permitted for any minute during the period. The <code>Maximum</code> statistic is the highest throughput permitted for any minute during the period. The <code>Average</code> statistic is the average throughput permitted during the period.</p> <p>Units: Bytes per second</p> <p>Valid statistics: <code>Minimum</code>, <code>Maximum</code>, <code>Average</code></p>
TotalIOBytes	<p>The number of bytes for each file system operation, including data read, data write, and metadata operations.</p> <p>The <code>Sum</code> statistic is the total number of bytes associated with all file system operations. The <code>Minimum</code> statistic is the size of the smallest operation during the period. The <code>Maximum</code> statistic is the size of the largest operation during the period. The <code>Average</code> statistic is the average size of an operation during the period. The <code>SampleCount</code> statistic provides a count of all operations.</p> <p>Note To calculate the average operations per second for a period, divide the <code>SampleCount</code> statistic by the number of seconds in the period. To calculate the average throughput (Bytes per second) for a period, divide the <code>Sum</code> statistic by the number of seconds in the period.</p> <p>Units:</p> <ul style="list-style-type: none"> • Bytes for <code>Minimum</code>, <code>Maximum</code>, <code>Average</code>, and <code>Sum</code> statistics. • Count for <code>SampleCount</code>. <p>Valid statistics: <code>Minimum</code>, <code>Maximum</code>, <code>Average</code>, <code>Sum</code>, <code>SampleCount</code></p>

Dimensions for Amazon EFS Metrics

Amazon EFS Dimensions

Amazon EFS metrics use the `EFS` namespace and provides metrics for a single dimension, `FileSystemId`. A file system's ID can be found in the Amazon EFS management console, and it takes the form of `fs-XXXXXXX`.

Elastic Load Balancing Metrics and Dimensions

Elastic Load Balancing supports two types of load balancers: Classic Load Balancers and Application Load Balancers. Elastic Load Balancing sends metrics to CloudWatch for both types of load balancers.

Contents

- [Application Load Balancer Metrics \(p. 85\)](#)

- [Metric Dimensions for Application Load Balancers \(p. 86\)](#)
- [Classic Load Balancer Metrics \(p. 86\)](#)
- [Metric Dimensions for Classic Load Balancers \(p. 89\)](#)

Application Load Balancer Metrics

The AWS/ApplicationELB namespace includes the following metrics.

Metric	Description
ActiveConnectionCount	The total number of concurrent TCP connections active from clients to the load balancer and from the load balancer to targets. Statistics: The most useful statistic is Sum.
ClientTLSNegotiationErrors	The number of TLS connections initiated by the client that did not establish a session with the load balancer. Possible causes include a mismatch of ciphers or protocols. Statistics: The most useful statistic is Sum.
HealthyHostCount	The number of targets that are considered healthy. Statistics: The most useful statistics are Average, Minimum, and Maximum.
HTTPCode_ELB_4XX_Count	The number of HTTP 4XX client error codes that originate from the load balancer. Client errors are generated when requests are malformed or incomplete. These requests have not been received by the target. This count does not include any response codes generated by the targets. Statistics: The most useful statistic is Sum. Note that Minimum, Maximum, and Average all return 1.
HTTPCode_ELB_5XX_Count	The number of HTTP 5XX server error codes that originate from the load balancer. This count does not include any response codes generated by the targets. Statistics: The most useful statistic is Sum. Note that Minimum, Maximum, and Average all return 1.
HTTPCode_Target_2XX_Count HTTPCode_Target_3XX_Count HTTPCode_Target_4XX_Count HTTPCode_Target_5XX_Count	The number of HTTP response codes generated by the targets. This does not include any response codes generated by the load balancer. Statistics: The most useful statistic is Sum. Note that Minimum, Maximum, and Average all return 1.
NewConnectionCount	The total number of new TCP connections established from clients to the load balancer and from the load balancer to targets. Statistics: The most useful statistic is Sum.
ProcessedBytes	The total number of bytes processed by the load balancer.
RejectedConnectionCount	The number of connections that were rejected because the load balancer could not establish a connection with a healthy target in order to route the request. Statistics: The most useful statistic is Sum.
RequestCount	The number of requests received by the load balancer. Statistics: The most useful statistic is Sum. Note that Minimum, Maximum, and Average all return 1.
TargetConnectionErrors	The number of connections that were not successfully established between the load balancer and target.

Metric	Description
	Statistics: The most useful statistic is <code>Sum</code> .
<code>TargetResponseTime</code>	The time elapsed, in seconds, after the request leaves the load balancer until a response from the target is received. This is equivalent to the <code>target_processing_time</code> field in the access logs. Statistics: The most useful statistics are <code>Average</code> and <code>pNN.NN</code> (percentiles).
<code>TargetTLSNegotiationErrors</code>	The number of TLS connections initiated by the load balancer that did not establish a session with the target. Possible causes include a mismatch of ciphers or protocols. Statistics: The most useful statistic is <code>Sum</code> .
<code>UnHealthyHostCount</code>	The number of targets that are considered unhealthy. Statistics: The most useful statistics are <code>Average</code> , <code>Minimum</code> , and <code>Maximum</code> .

Metric Dimensions for Application Load Balancers

To filter the metrics for your Application Load Balancer, use the following dimensions.

Dimension	Description
<code>AvailabilityZone</code>	Filter the metric data by Availability Zone.
<code>LoadBalancer</code>	Filter the metric data by load balancer. Specify the load balancer as follows: <code>app/load-balancer-name/1234567890123456</code> (the final portion of the load balancer ARN).
<code>TargetGroup</code>	Filter the metric data by target group. Specify the target group as follows: <code>targetgroup/target-group-name/1234567890123456</code> (the final portion of the target group ARN).

Classic Load Balancer Metrics

The `AWS/ELB` namespace includes the following metrics.

Metric	Description
<code>BackendConnectionErrors</code>	The number of connections that were not successfully established between the load balancer and the registered instances. Because the load balancer retries the connection when there are errors, this count can exceed the request rate. Note that this count also includes any connection errors related to health checks. Reporting criteria: There is a nonzero value Statistics: The most useful statistic is <code>Sum</code> . Note that <code>Average</code> , <code>Minimum</code> , and <code>Maximum</code> are reported per load balancer node and are not typically useful. However, the difference between the minimum and maximum (or peak to average or average to trough) might be useful to determine whether a load balancer node is an outlier. Example: Suppose that your load balancer has 2 instances in <code>us-west-2a</code> and 2 instances in <code>us-west-2b</code> , and that attempts to connect to 1 instance in <code>us-west-2a</code> result in back-end connection errors. The sum for <code>us-west-2a</code> includes these connection errors, while the sum for <code>us-west-2b</code> does not

Metric	Description
	include them. Therefore, the sum for the load balancer equals the sum for us-west-2a.
HealthyHostCount, UnHealthyHostCount	<p>The number of healthy and unhealthy instances registered with your load balancer. A newly registered instance is considered healthy after it passes the first health check. An instance is considered unhealthy after it exceeds the unhealthy threshold configured for health checks. An unhealthy instance is considered healthy again after it meets the healthy threshold configured for health checks. If cross-zone load balancing is enabled, the number of healthy instances for the <code>LoadBalancerName</code> dimension is calculated across all Availability Zones.</p> <p>Reporting criteria: There are registered instances</p> <p>Statistics: The most useful statistics are <code>Average</code>, <code>Minimum</code>, and <code>Maximum</code>. These statistics are determined by the load balancer nodes. Note that some load balancer nodes might determine that an instance is unhealthy for a brief period while other nodes determine that it is healthy.</p> <p>Example: Suppose that your load balancer has 2 instances in us-west-2a and 2 instances in us-west-2b, us-west-2a has 1 unhealthy instance, and us-west-2b has no unhealthy instances. With the <code>AvailabilityZone</code> dimension, there is an average of 1 healthy and 1 unhealthy instance in us-west-2a, and an average of 2 healthy and 0 unhealthy instances in us-west-2b.</p>
HTTPCode_Backend_2XX, HTTPCode_Backend_3XX, HTTPCode_Backend_4XX, HTTPCode_Backend_5XX	<p>[HTTP listener] The number of HTTP response codes generated by registered instances. This count does not include any response codes generated by the load balancer.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is <code>Sum</code>. Note that <code>Minimum</code>, <code>Maximum</code>, and <code>Average</code> are all 1.</p> <p>Example: Suppose that your load balancer has 2 instances in us-west-2a and 2 instances in us-west-2b, and that requests sent to 1 instance in us-west-2a result in HTTP 500 responses. The sum for us-west-2a includes these error responses, while the sum for us-west-2b does not include them. Therefore, the sum for the load balancer equals the sum for us-west-2a.</p>
HTTPCode_ELB_4XX	<p>[HTTP listener] The number of HTTP 4XX client error codes generated by the load balancer. Client errors are generated when a request is malformed or incomplete.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is <code>Sum</code>. Note that <code>Minimum</code>, <code>Maximum</code>, and <code>Average</code> are all 1.</p> <p>Example: Suppose that your load balancer has us-west-2a and us-west-2b enabled, and that client requests include a malformed request URL. As a result, client errors would likely increase in all Availability Zones. The sum for the load balancer is the sum of the values for the Availability Zones.</p>

Metric	Description
HTTPCode_ELB_5XX	<p>[HTTP listener] The number of HTTP 5XX server error codes generated by the load balancer. This count does not include any response codes generated by the registered instances. The metric is reported if there are no healthy instances registered to the load balancer, or if the request rate exceeds the capacity of the instances (spillover) or the load balancer.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is <code>Sum</code>. Note that <code>Minimum</code>, <code>Maximum</code>, and <code>Average</code> are all 1.</p> <p>Example: Suppose that your load balancer has <code>us-west-2a</code> and <code>us-west-2b</code> enabled, and that instances in <code>us-west-2a</code> are experiencing high latency and are slow to respond to requests. As a result, the surge queue for the load balancer nodes in <code>us-west-2a</code> fills and clients receive a 503 error. If <code>us-west-2b</code> continues to respond normally, the sum for the load balancer equals the sum for <code>us-west-2a</code>.</p>
Latency	<p>[HTTP listener] The time elapsed, in seconds, after the request leaves the load balancer until the headers of the response are received.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistics are <code>Average</code> and <code>pNN.NN</code> (percentiles). Use <code>Maximum</code> to determine whether some requests are taking substantially longer than the average. Note that <code>Minimum</code> is typically not useful.</p> <p>Example: Suppose that your load balancer has 2 instances in <code>us-west-2a</code> and 2 instances in <code>us-west-2b</code>, and that requests sent to 1 instance in <code>us-west-2a</code> have a higher latency. The average for <code>us-west-2a</code> has a higher value than the average for <code>us-west-2b</code>.</p>
RequestCount	<p>The number of requests completed or connections made during the specified interval (1 or 5 minutes).</p> <p>[HTTP listener] The number of requests received and routed, including HTTP error responses from the registered instances.</p> <p>[TCP listener] The number of connections made to the registered instances.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is <code>Sum</code>. Note that <code>Minimum</code>, <code>Maximum</code>, and <code>Average</code> all return 1.</p> <p>Example: Suppose that your load balancer has 2 instances in <code>us-west-2a</code> and 2 instances in <code>us-west-2b</code>, and that 100 requests are sent to the load balancer. There are 60 requests sent to <code>us-west-2a</code>, with each instance receiving 30 requests, and 40 requests sent to <code>us-west-2b</code>, with each instance receiving 20 requests. With the <code>AvailabilityZone</code> dimension, there is a sum of 60 requests in <code>us-west-2a</code> and 40 requests in <code>us-west-2b</code>. With the <code>LoadBalancerName</code> dimension, there is a sum of 100 requests.</p>

Metric	Description
SpilloverCount	<p>The total number of requests that were rejected because the surge queue is full.</p> <p>[HTTP listener] The load balancer returns an HTTP 503 error code. [TCP listener] The load balancer closes the connection.</p> <p>Reporting criteria: There is a nonzero value</p> <p>Statistics: The most useful statistic is <code>Sum</code>. Note that <code>Average</code>, <code>Minimum</code>, and <code>Maximum</code> are reported per load balancer node and are not typically useful.</p> <p>Example: Suppose that your load balancer has us-west-2a and us-west-2b enabled, and that instances in us-west-2a are experiencing high latency and are slow to respond to requests. As a result, the surge queue for the load balancer node in us-west-2a fills, resulting in spillover. If us-west-2b continues to respond normally, the sum for the load balancer will be the same as the sum for us-west-2a.</p>
SurgeQueueLength	<p>The total number of requests that are pending routing. The load balancer queues a request if it is unable to establish a connection with a healthy instance in order to route the request. The maximum size of the queue is 1,024. Additional requests are rejected when the queue is full. For more information, see <code>SpilloverCount</code>.</p> <p>Reporting criteria: There is a nonzero value.</p> <p>Statistics: The most useful statistic is <code>Maximum</code>, because it represents the peak of queued requests. The <code>Average</code> statistic can be useful in combination with <code>Minimum</code> and <code>Maximum</code> to determine the range of queued requests. Note that <code>Sum</code> is not useful.</p> <p>Example: Suppose that your load balancer has us-west-2a and us-west-2b enabled, and that instances in us-west-2a are experiencing high latency and are slow to respond to requests. As a result, the surge queue for the load balancer nodes in us-west-2a fills, with clients likely experiencing increased response times. If this continues, the load balancer will likely have spillovers (see the <code>SpilloverCount</code> metric). If us-west-2b continues to respond normally, the <code>max</code> for the load balancer will be the same as the <code>max</code> for us-west-2a.</p>

Metric Dimensions for Classic Load Balancers

To filter the metrics for your Classic Load Balancer, use the following dimensions.

Dimension	Description
AvailabilityZone	Filter the metric data by the specified Availability Zone.
LoadBalancerName	Filter the metric data by the specified load balancer.

Amazon EMR Metrics and Dimensions

Amazon EMR (Amazon EMR) sends metrics to CloudWatch. All Amazon EMR job flows automatically send metrics in five-minute intervals. Metrics are archived for two weeks; after that period, the data is discarded. For more information, see [Monitor Metrics with Amazon CloudWatch](#) in the *Amazon EMR Developer Guide*.

Amazon EMR Metrics

Amazon EMR sends the following metrics to Amazon CloudWatch.

The `AWS/ElasticMapReduce` namespace includes the following metrics.

Note

Amazon EMR pulls metrics from a cluster. If a cluster becomes unreachable, no metrics are reported until the cluster becomes available again.

The following are Hadoop 1 metrics:

Metric	Description
<i>Cluster Status</i>	
IsIdle	Indicates that a cluster is no longer performing work, but is still alive and accruing charges. It is set to 1 if no tasks are running and no jobs are running, and set to 0 otherwise. This value is checked at five-minute intervals and a value of 1 indicates only that the cluster was idle when checked, not that it was idle for the entire five minutes. To avoid false positives, you should raise an alarm when this value has been 1 for more than one consecutive 5-minute check. For example, you might raise an alarm on this value if it has been 1 for thirty minutes or longer. Use case: Monitor cluster performance Units: <i>Boolean</i>
JobsRunning	The number of jobs in the cluster that are currently running. Use case: Monitor cluster health Units: <i>Count</i>
JobsFailed	The number of jobs in the cluster that have failed. Use case: Monitor cluster health Units: <i>Count</i>
<i>Map/Reduce</i>	
MapTasksRunning	The number of running map tasks for each job. If you have a scheduler installed and multiple jobs running, multiple graphs are generated. Use case: Monitor cluster progress Units: <i>Count</i>
MapTasksRemaining	The number of remaining map tasks for each job. If you have a scheduler installed and multiple jobs running, multiple graphs are generated. A remaining map task is one that is not in any of the following states: Running, Killed, or Completed. Use case: Monitor cluster progress Units: <i>Count</i>

Metric	Description
MapSlotsOpen	<p>The unused map task capacity. This is calculated as the maximum number of map tasks for a given cluster, less the total number of map tasks currently running in that cluster.</p> <p>Use case: Analyze cluster performance</p> <p>Units: <i>Count</i></p>
RemainingMapTasksPerSlot	<p>The ratio of the total map tasks remaining to the total map slots available in the cluster.</p> <p>Use case: Analyze cluster performance</p> <p>Units: <i>Ratio</i></p>
ReduceTasksRunning	<p>The number of running reduce tasks for each job. If you have a scheduler installed and multiple jobs running, multiple graphs are generated.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Count</i></p>
ReduceTasksRemaining	<p>The number of remaining reduce tasks for each job. If you have a scheduler installed and multiple jobs running, multiple graphs are generated.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Count</i></p>
ReduceSlotsOpen	<p>Unused reduce task capacity. This is calculated as the maximum reduce task capacity for a given cluster, less the number of reduce tasks currently running in that cluster.</p> <p>Use case: Analyze cluster performance</p> <p>Units: <i>Count</i></p>
<i>Node Status</i>	
CoreNodesRunning	<p>The number of core nodes working. Data points for this metric are reported only when a corresponding instance group exists.</p> <p>Use case: Monitor cluster health</p> <p>Units: <i>Count</i></p>
CoreNodesPending	<p>The number of core nodes waiting to be assigned. All of the core nodes requested may not be immediately available; this metric reports the pending requests. Data points for this metric are reported only when a corresponding instance group exists.</p> <p>Use case: Monitor cluster health</p> <p>Units: <i>Count</i></p>

Metric	Description
LiveDataNodes	<p>The percentage of data nodes that are receiving work from Hadoop.</p> <p>Use case: Monitor cluster health</p> <p>Units: <i>Percent</i></p>
TaskNodesRunning	<p>The number of task nodes working. Data points for this metric are reported only when a corresponding instance group exists.</p> <p>Use case: Monitor cluster health</p> <p>Units: <i>Count</i></p>
TaskNodesPending	<p>The number of core nodes waiting to be assigned. All of the task nodes requested may not be immediately available; this metric reports the pending requests. Data points for this metric are reported only when a corresponding instance group exists.</p> <p>Use case: Monitor cluster health</p> <p>Units: <i>Count</i></p>
LiveTaskTrackers	<p>The percentage of task trackers that are functional.</p> <p>Use case: Monitor cluster health</p> <p>Units: <i>Percent</i></p>
<i>IO</i>	
S3BytesWritten	<p>The number of bytes written to Amazon S3.</p> <p>Use case: Analyze cluster performance, Monitor cluster progress</p> <p>Units: <i>Bytes</i></p>
S3BytesRead	<p>The number of bytes read from Amazon S3.</p> <p>Use case: Analyze cluster performance, Monitor cluster progress</p> <p>Units: <i>Bytes</i></p>
HDFSUtilization	<p>The percentage of HDFS storage currently used.</p> <p>Use case: Analyze cluster performance</p> <p>Units: <i>Percent</i></p>
HDFSBytesRead	<p>The number of bytes read from HDFS.</p> <p>Use case: Analyze cluster performance, Monitor cluster progress</p> <p>Units: <i>Bytes</i></p>

Metric	Description
HDFSBytesWritten	The number of bytes written to HDFS. Use case: Analyze cluster performance, Monitor cluster progress Units: <i>Bytes</i>
MissingBlocks	The number of blocks in which HDFS has no replicas. These might be corrupt blocks. Use case: Monitor cluster health Units: <i>Count</i>
TotalLoad	The total number of concurrent data transfers. Use case: Monitor cluster health Units: <i>Count</i>
<i>HBase</i>	
BackupFailed	Whether the last backup failed. This is set to 0 by default and updated to 1 if the previous backup attempt failed. This metric is only reported for HBase clusters. Use case: Monitor HBase backups Units: <i>Count</i>
MostRecentBackupDuration	The amount of time it took the previous backup to complete. This metric is set regardless of whether the last completed backup succeeded or failed. While the backup is ongoing, this metric returns the number of minutes after the backup started. This metric is only reported for HBase clusters. Use case: Monitor HBase Backups Units: <i>Minutes</i>
TimeSinceLastSuccessfulBackup	The number of elapsed minutes after the last successful HBase backup started on your cluster. This metric is only reported for HBase clusters. Use case: Monitor HBase backups Units: <i>Minutes</i>

The following metrics are available for Hadoop 2 AMIs:

Metric	Description
<i>Cluster Status</i>	

Metric	Description
IsIdle	<p>Indicates that a cluster is no longer performing work, but is still alive and accruing charges. It is set to 1 if no tasks are running and no jobs are running, and set to 0 otherwise. This value is checked at five-minute intervals and a value of 1 indicates only that the cluster was idle when checked, not that it was idle for the entire five minutes. To avoid false positives, you should raise an alarm when this value has been 1 for more than one consecutive 5-minute check. For example, you might raise an alarm on this value if it has been 1 for thirty minutes or longer.</p> <p>Use case: Monitor cluster performance</p> <p>Units: <i>Boolean</i></p>
ContainerAllocated	<p>The number of resource containers allocated by the ResourceManager.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Count</i></p>
ContainerReserved	<p>The number of containers reserved.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Count</i></p>
ContainerPending	<p>The number of containers in the queue that have not yet been allocated.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Count</i></p>
AppsCompleted	<p>The number of applications submitted to YARN that have completed.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Count</i></p>
AppsFailed	<p>The number of applications submitted to YARN that have failed to complete.</p> <p>Use case: Monitor cluster progress, Monitor cluster health</p> <p>Units: <i>Count</i></p>
AppsKilled	<p>The number of applications submitted to YARN that have been killed.</p> <p>Use case: Monitor cluster progress, Monitor cluster health</p> <p>Units: <i>Count</i></p>

Metric	Description
AppsPending	<p>The number of applications submitted to YARN that are in a pending state.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Count</i></p>
AppsRunning	<p>The number of applications submitted to YARN that are running.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Count</i></p>
AppsSubmitted	<p>The number of applications submitted to YARN.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Count</i></p>
<i>Node Status</i>	
CoreNodesRunning	<p>The number of core nodes working. Data points for this metric are reported only when a corresponding instance group exists.</p> <p>Use case: Monitor cluster health</p> <p>Units: <i>Count</i></p>
CoreNodesPending	<p>The number of core nodes waiting to be assigned. All of the core nodes requested may not be immediately available; this metric reports the pending requests. Data points for this metric are reported only when a corresponding instance group exists.</p> <p>Use case: Monitor cluster health</p> <p>Units: <i>Count</i></p>
LiveDataNodes	<p>The percentage of data nodes that are receiving work from Hadoop.</p> <p>Use case: Monitor cluster health</p> <p>Units: <i>Percent</i></p>
MRTotalNodes	<p>The number of nodes presently available to MapReduce jobs.</p> <p>Use ase: Monitor cluster progress</p> <p>Units: <i>Count</i></p>
MRActiveNodes	<p>The number of nodes presently running MapReduce tasks or jobs.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Count</i></p>

Metric	Description
MRLostNodes	<p>The number of nodes allocated to MapReduce that have been marked in a LOST state.</p> <p>Use case: Monitor cluster health, Monitor cluster progress</p> <p>Units: <i>Count</i></p>
MRUnhealthyNodes	<p>The number of nodes available to MapReduce jobs marked in an UNHEALTHY state.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Count</i></p>
MRDecommissionedNodes	<p>The number of nodes allocated to MapReduce applications that have been marked in a DECOMMISSIONED state.</p> <p>Use ase: Monitor cluster health, Monitor cluster progress</p> <p>Units: <i>Count</i></p>
MRRebootedNodes	<p>The number of nodes available to MapReduce that have been rebooted and marked in a REBOOTED state.</p> <p>Use case: Monitor cluster health, Monitor cluster progress</p> <p>Units: <i>Count</i></p>
<i>IO</i>	
S3BytesWritten	<p>The number of bytes written to Amazon S3.</p> <p>Use case: Analyze cluster performance, Monitor cluster progress</p> <p>Units: <i>Bytes</i></p>
S3BytesRead	<p>The number of bytes read from Amazon S3.</p> <p>Use case: Analyze cluster performance, Monitor cluster progress</p> <p>Units: <i>Bytes</i></p>
HDFSUtilization	<p>The percentage of HDFS storage currently used.</p> <p>Use case: Analyze cluster performance</p> <p>Units: <i>Percent</i></p>
HDFSBytesRead	<p>The number of bytes read from HDFS.</p> <p>Use case: Analyze cluster performance, Monitor cluster progress</p> <p>Units: <i>Bytes</i></p>

Metric	Description
HDFSBytesWritten	<p>The number of bytes written to HDFS.</p> <p>Use case: Analyze cluster performance, Monitor cluster progress</p> <p>Units: <i>Bytes</i></p>
MissingBlocks	<p>The number of blocks in which HDFS has no replicas. These might be corrupt blocks.</p> <p>Use case: Monitor cluster health</p> <p>Units: <i>Count</i></p>
CorruptBlocks	<p>The number of blocks that HDFS reports as corrupted.</p> <p>Use case: Monitor cluster health</p> <p>Units: <i>Count</i></p>
TotalLoad	<p>The total number of concurrent data transfers.</p> <p>Use case: Monitor cluster health</p> <p>Units: <i>Count</i></p>
MemoryTotalMB	<p>The total amount of memory in the cluster.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Bytes</i></p>
MemoryReservedMB	<p>The amount of memory reserved.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Bytes</i></p>
MemoryAvailableMB	<p>The amount of memory available to be allocated.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Bytes</i></p>
MemoryAllocatedMB	<p>The amount of memory allocated to the cluster.</p> <p>Use case: Monitor cluster progress</p> <p>Units: <i>Bytes</i></p>
PendingDeletionBlocks	<p>The number of blocks marked for deletion.</p> <p>Use case: Monitor cluster progress, Monitor cluster health</p> <p>Units: <i>Count</i></p>

Metric	Description
UnderReplicatedBlocks	<p>The number of blocks that need to be replicated one or more times.</p> <p>Use case: Monitor cluster progress, Monitor cluster health</p> <p>Units: <i>Count</i></p>
DfsPendingReplicationBlocks	<p>The status of block replication: blocks being replicated, age of replication requests, and unsuccessful replication requests.</p> <p>Use case: Monitor cluster progress, Monitor cluster health</p> <p>Units: <i>Count</i></p>
CapacityRemainingGB	<p>The amount of remaining HDFS disk capacity.</p> <p>Use case: Monitor cluster progress, Monitor cluster health</p> <p>Units: <i>Bytes</i></p>
<i>HBase</i>	
HbaseBackupFailed	<p>Whether the last backup failed. This is set to 0 by default and updated to 1 if the previous backup attempt failed. This metric is only reported for HBase clusters.</p> <p>Use case: Monitor HBase backups</p> <p>Units: <i>Count</i></p>
MostRecentBackupDuration	<p>The amount of time it took the previous backup to complete. This metric is set regardless of whether the last completed backup succeeded or failed. While the backup is ongoing, this metric returns the number of minutes after the backup started. This metric is only reported for HBase clusters.</p> <p>Use case: Monitor HBase Backups</p> <p>Units: <i>Minutes</i></p>
TimeSinceLastSuccessfulBackup	<p>The number of elapsed minutes after the last successful HBase backup started on your cluster. This metric is only reported for HBase clusters.</p> <p>Use case: Monitor HBase backups</p> <p>Units: <i>Minutes</i></p>

Amazon EMR Dimensions

The following dimensions are available for Amazon EMR.

Dimension	Description
ClusterId/JobFlowId	<p>The identifier for a cluster. You can find this value by clicking on the cluster in the Amazon EMR console. It takes the form <code>j-xxxxxxxxxxxxxx</code>.</p>

Dimension	Description
JobId	The identifier of a job within a cluster. You can use this to filter the metrics returned from a cluster down to those that apply to a single job within the cluster. JobId takes the form <code>job_XXXXXXXXXXXXX_XXXX</code> .

Amazon Elasticsearch Service Metrics and Dimensions

Amazon Elasticsearch Service sends data to CloudWatch every minute. You can create alarms using [Amazon Elasticsearch Service Metrics and Dimensions \(p. 99\)](#). For more information, see [Monitoring Cluster Metrics and Statistics with Amazon CloudWatch](#) in the *Amazon Elasticsearch Service Developer Guide*.

Amazon Elasticsearch Service Metrics

The `AWS/ES` namespace includes the following metrics for clusters.

Metric	Description
ClusterStatus.green	Indicates that all index shards are allocated to nodes in the cluster. Relevant statistics: Minimum, Maximum
ClusterStatus.yellow	Indicates that the primary shards for all indices are allocated to nodes in a cluster, but the replica shards for at least one index are not. Note that single node clusters always initialize with this cluster status because there is no second node to which a replica can be assigned. You can either increase your node count to obtain a green cluster status, or you can use the Amazon ES API to set the <code>number_of_replicas</code> setting for your index to 0. For more information, see Update Indices Settings in the Amazon ES documentation. Relevant statistics: Minimum, Maximum
ClusterStatus.red	Indicates that the primary and replica shards of at least one index are not allocated to nodes in a cluster. For more information, see Red Cluster Status . Relevant statistics: Minimum, Maximum
Nodes	The number of nodes in the Amazon ES cluster. Relevant Statistics: Minimum, Maximum, Average
SearchableDocuments	The total number of searchable documents across all indices in the cluster. Relevant statistics: Minimum, Maximum, Average
DeletedDocuments	The total number of deleted documents across all indices in the cluster. Relevant statistics: Minimum, Maximum, Average
CPUUtilization	The maximum percentage of CPU resources used for data nodes in the cluster.

Metric	Description
	Relevant statistics: Maximum, Average
FreeStorageSpace	The free space, in megabytes, for all data nodes in the cluster. Relevant statistics: Minimum
ClusterUsedSpace	The total used space, in megabytes, for a cluster. You can view this metric in the Amazon CloudWatch console, but not in the Amazon ES console. Relevant statistics: Minimum, Maximum
JVMMemoryPressure	The maximum percentage of the Java heap used for all data nodes in the cluster. Relevant statistics: Maximum
AutomatedSnapshotFailure	The number of failed automated snapshots for the cluster. A value of 1 indicates that no automated snapshot was taken for the domain in the previous 36 hours. Relevant statistics: Minimum, Maximum
CPUCreditBalance	The remaining CPU credits available for data nodes in the cluster. A CPU credit provides the performance of a full CPU core for one minute. This metrics is available only for the t2.micro.elasticsearch, t2.small.elasticsearch, and t2.medium.elasticsearch instance types. Relevant statistics: Minimum

The `AWS/ES` namespace includes the following metrics for dedicated master nodes.

Metric	Description
MasterCPUUtilization	The maximum percentage of CPU resources used by the dedicated master nodes. We recommend increasing the size of the instance type when this metric reaches 60 percent. Relevant statistics: Average
MasterFreeStorageSpace	This metric is not relevant and can be ignored. The service does not use master nodes as data nodes.
MasterJVMMemoryPressure	The maximum percentage of the Java heap used for all dedicated master nodes in the cluster. We recommend moving to a larger instance type when this metric reaches 85 percent. Relevant statistics: Maximum
MasterCPUCreditBalance	The remaining CPU credits available for dedicated master nodes in the cluster. A CPU credit provides the performance of a full CPU core for one minute. This metric is available only for the t2.micro.elasticsearch, t2.small.elasticsearch, and t2.medium.elasticsearch instance types. Relevant statistics: Minimum

The `AWS/EBS` namespace includes the following metrics for EBS volumes.

Metric	Description
ReadLatency	The latency, in seconds, for read operations on EBS volumes. Relevant statistics: Minimum, Maximum, Average
WriteLatency	The latency, in seconds, for write operations on EBS volumes. Relevant statistics: Minimum, Maximum, Average
ReadThroughput	The throughput, in bytes per second, for read operations on EBS volumes. Relevant statistics: Minimum, Maximum, Average
WriteThroughput	The throughput, in bytes per second, for write operations on EBS volumes. Relevant statistics: Minimum, Maximum, Average
DiskQueueDepth	The number of pending input and output (I/O) requests for an EBS volume. Relevant statistics: Minimum, Maximum, Average
ReadIOPS	The number of input and output (I/O) operations per second for read operations on EBS volumes. Relevant statistics: Minimum, Maximum, Average
WriteIOPS	The number of input and output (I/O) operations per second for write operations on EBS volumes. Relevant statistics: Minimum, Maximum, Average

Dimensions for Amazon Elasticsearch Service Metrics

To filter the metrics, use the following dimensions.

Dimension	Description
ClientId	The AWS account ID.
DomainName	The name of the search domain.

Amazon Elastic Transcoder Metrics and Dimensions

When you interact with Amazon Elastic Transcoder, it sends the following metrics to CloudWatch every minute.

Elastic Transcoder Metrics

The `AWS/ElasticTranscoder` namespace includes the following metrics.

Metric	Description
Billed HD Output	<p>The number of billable seconds of HD output for a pipeline.</p> <p>Valid Dimensions: PipelineId</p> <p>Unit: Seconds</p>
Billed SD Output	<p>The number of billable seconds of SD output for a pipeline.</p> <p>Valid Dimensions: PipelineId</p> <p>Unit: Seconds</p>
Billed Audio Output	<p>The number of billable seconds of audio output for a pipeline.</p> <p>Valid Dimensions: PipelineId</p> <p>Unit: Seconds</p>
Jobs Completed	<p>The number of jobs completed by this pipeline.</p> <p>Valid Dimensions: PipelineId</p> <p>Unit: Count</p>
Jobs Errored	<p>The number of jobs that failed because of invalid inputs, such as a request to transcode a file that is not in the given input bucket.</p> <p>Valid Dimensions: PipelineId</p> <p>Unit: Count</p>
Outputs per Job	<p>The number of outputs Elastic Transcoder created for a job.</p> <p>Valid Dimensions: PipelineId</p> <p>Unit: Count</p>
Standby Time	<p>The number of seconds before Elastic Transcoder started transcoding a job.</p> <p>Valid Dimensions: PipelineId</p> <p>Unit: Seconds</p>
Errors	<p>The number of errors caused by invalid operation parameters, such as a request for a job status that does not include the job ID.</p> <p>Valid Dimensions: Operation</p> <p>Unit: Count</p>

Metric	Description
Throttles	The number of times that Elastic Transcoder automatically throttled an operation. Valid Dimensions: Operation Unit: Count

Dimensions for Elastic Transcoder Metrics

Elastic Transcoder metrics use the Elastic Transcoder namespace and provide metrics for the following dimension(s):

Dimension	Description
PipelineId	The ID of a pipeline. This dimension filters the data you request for an Elastic Transcoder pipeline.
Operation	This dimension filters the data you request for the APIs that Elastic Transcoder provides.

AWS IoT Metrics and Dimensions

When you interact with AWS IoT, it sends the following metrics to CloudWatch every minute.

AWS IoT Metrics

The `AWS/IoT` namespace includes the following metrics.

AWS IoT sends the following metrics to CloudWatch once per received request.

Metric	Description
PublishIn.Success	A client published on an MQTT topic successfully. Valid Dimensions: Protocol Valid Statistics:1 for success, 0 for failure. Unit: Count
PublishOut.Success	Clients subscribed to an MQTT topic recieved a published message. Valid Dimensions: Protocol Valid Statistics:1 for success, 0 for failure. Unit: Count
Subscribe.Success	AWS IoT message broker received a request to subscribe to an MQTT topic. Valid Dimensions: Protocol Valid Statistics:1 for success, 0 for failure.

Metric	Description
	Unit: Count
Ping.Success	AWS IoT received a Ping message. Valid Dimensions: Protocol Valid Statistics:1 per ping request from the client. Unit: Count
Connect.Success	A client connected to AWS IoT. Valid Dimensions: Protocol Valid Statistics: 1 per successful MQTT connection from the client. Unit: Count
GetThingShadow.Accepted	AWS IoT received a GetThingShadow request. Valid Dimensions: Protocol Valid Statistics:1 for success, 0 for failure. Unit: Count
UpdateThingShadow.Accepted	AWS IoT received a UpdateThingShadow request. Valid Dimensions: Protocol Valid Statistics:1 for success, 0 for failure. Unit: Count
DeleteThingShadow.Accepted	AWS IoT received a DeleteThingShadow request. Valid Dimensions: Protocol Valid Statistics:1 for success, 0 for failure. Unit: Count
RulesExecuted	AWS IoT executed a rule.. Valid Dimensions: Protocol Valid Statistics:1 for success, 0 for failure. Unit: Count

Dimensions for AWS IoT Metrics

Metrics use the namespace and provide metrics for the following dimension(s):

Dimension	Description
Protocol	The protocol with which the request was made. Valid values are MQTT or HTTP.

Amazon Kinesis Analytics Metrics

Analytics sends metrics to CloudWatch. For more information, see [Monitoring with Amazon CloudWatch Metrics](#) in the *Amazon Kinesis Analytics Developer Guide*.

Metrics

The `AWS/KinesisAnalytics` namespace includes the following metrics.

Metric	Description
Bytes	The number of bytes read (per input stream) or written (per output stream). Levels: Per input stream and per output stream
MillisBehindLatest	Indicates how far behind from the current time an application is reading from the streaming source. Levels: Application-level
Records	The number of records read (per input stream) or written (per output stream). Levels: Per input stream and per output stream

Dimensions for Metrics

Amazon Kinesis Analytics provides metrics for the following dimensions.

Dimension	Description
Flow	Per input stream: Input Per output stream: Output
Id	Per input stream: Input Id Per output stream: Output Id

Amazon Kinesis Firehose Metrics

Firehose sends metrics to CloudWatch. For more information, see [Monitoring with Amazon CloudWatch Metrics](#) in the *Amazon Kinesis Firehose Developer Guide*.

Service-level CloudWatch Metrics

The `AWS/Firehose` namespace includes the following service-level metrics.

Metric	Description
DeliveryToElasticsearch	The number of bytes indexed to Amazon ES over the specified time period. Units: Bytes

Metric	Description
DeliveryToElasticsearch	The number of records indexed to Amazon ES over the specified time period. Units: Count
DeliveryToElasticsearch.Successful	The sum of the successfully indexed records over the sum of records that were attempted.
DeliveryToRedshift.Bytes	The number of bytes copied to Amazon Redshift over the specified time period. Units: Bytes
DeliveryToRedshift.Records	The number of records copied to Amazon Redshift over the specified time period. Units: Count
DeliveryToRedshift.Successful	The sum of successful Amazon Redshift COPY commands over the sum of all Amazon Redshift COPY commands.
DeliveryToS3.Bytes	The number of bytes delivered to Amazon S3 over the specified time period. Units: Bytes
DeliveryToS3.DataFreshness	The age (from getting into Firehose to now) of the oldest record in Firehose. Any record older than this age has been delivered to the S3 bucket. Units: Seconds
DeliveryToS3.Records	The number of records delivered to Amazon S3 over the specified time period. Units: Count
DeliveryToS3.Successful	The sum of successful Amazon S3 put commands over the sum of all Amazon S3 put commands.
IncomingBytes	The number of bytes ingested into the Firehose stream over the specified time period. Units: Bytes
IncomingRecords	The number of records ingested into the Firehose stream over the specified time period. Units: Count

API-Level CloudWatch Metrics

The `AWS/Firehose` namespace includes the following API-level metrics.

Metric	Description
DescribeDeliveryStream	The time taken per <code>DescribeDeliveryStream</code> operation, measured over the specified time period.

Metric	Description
	Units: Milliseconds
DescribeDeliveryStreams	The total number of DescribeDeliveryStream requests. Units: Count
ListDeliveryStreams	The time taken per ListDeliveryStream operation, measured over the specified time period. Units: Milliseconds
ListDeliveryStreams	The total number of ListFirehose requests. Units: Count
PutRecord.Bytes	The number of bytes put to the Firehose delivery stream using PutRecord over the specified time period. Units: Bytes
PutRecord.Latency	The time taken per PutRecord operation, measured over the specified time period. Units: Milliseconds
PutRecord.Requests	The total number of PutRecord requests, which is equal to total number of records from PutRecord operations. Units: Count
PutRecordBatch.Bytes	The number of bytes put to the Firehose delivery stream using PutRecordBatch over the specified time period. Units: Bytes
PutRecordBatch.Latency	The time taken per PutRecordBatch operation, measured over the specified time period. Units: Milliseconds
PutRecordBatch.Records	The total number of records from PutRecordBatch operations. Units: Count
PutRecordBatch.Requests	The total number of PutRecordBatch requests. Units: Count
UpdateDeliveryStream	The time taken per UpdateDeliveryStream operation, measured over the specified time period. Units: Milliseconds
UpdateDeliveryStream	The total number of UpdateDeliveryStream requests. Units: Count

Amazon Kinesis Streams Metrics and Dimensions

Streams sends metrics to CloudWatch at two levels; the stream level and, optionally, the shard level. Stream-level metrics are for most common monitoring use cases in normal conditions. Shard-level metrics are for specific monitoring tasks, usually related to troubleshooting. For more information, see [Monitoring Amazon Kinesis with Amazon CloudWatch](#) in the *Amazon Kinesis Developer Guide*.

Contents

- [Basic Stream-level Metrics \(p. 108\)](#)
- [Enhanced Shard-level Metrics \(p. 111\)](#)
- [Dimensions for Amazon Kinesis Metrics \(p. 113\)](#)

Basic Stream-level Metrics

The `AWS/Kinesis` namespace includes the following stream-level metrics.

Streams sends these stream-level metrics to CloudWatch every minute. These metrics are always available.

Metric	Description
<code>GetRecords.Bytes</code>	<p>The number of bytes retrieved from the Amazon Kinesis stream, measured over the specified time period. Minimum, Maximum, and Average statistics represent the bytes in a single <code>GetRecords</code> operation for the stream in the specified time period.</p> <p>Shard-level metric name: <code>OutgoingBytes</code></p> <p>Dimensions: <code>StreamName</code></p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Bytes</p>
<code>GetRecords.IteratorAge</code>	<p>This metric is deprecated. Use <code>GetRecords.IteratorAgeMilliseconds</code>.</p>
<code>GetRecords.IteratorAgeMilliseconds</code>	<p>The age of the last record in all <code>GetRecords</code> calls made against an Amazon Kinesis stream, measured over the specified time period. Age is the difference between the current time and when the last record of the <code>GetRecords</code> call was written to the stream. The Minimum and Maximum statistics can be used to track the progress of Amazon Kinesis consumer applications. A value of zero indicates that the records being read are completely caught up with the stream.</p> <p>Shard-level metric name: <code>IteratorAgeMilliseconds</code></p> <p>Dimensions: <code>StreamName</code></p> <p>Statistics: Minimum, Maximum, Average, Samples</p> <p>Units: Milliseconds</p>
<code>GetRecords.Latency</code>	<p>The time taken per <code>GetRecords</code> operation, measured over the specified time period.</p> <p>Dimensions: <code>StreamName</code></p>

Metric	Description
	<p>Statistics: Minimum, Maximum, Average</p> <p>Units: Milliseconds</p>
GetRecords.Records	<p>The number of records retrieved from the shard, measured over the specified time period. Minimum, Maximum, and Average statistics represent the records in a single <code>GetRecords</code> operation for the stream in the specified time period.</p> <p>Shard-level metric name: <code>OutgoingRecords</code></p> <p>Dimensions: <code>StreamName</code></p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Count</p>
GetRecords.Success	<p>The number of successful <code>GetRecords</code> operations per stream, measured over the specified time period.</p> <p>Dimensions: <code>StreamName</code></p> <p>Statistics: Average, Sum, Samples</p> <p>Units: Count</p>
IncomingBytes	<p>The number of bytes successfully put to the Amazon Kinesis stream over the specified time period. This metric includes bytes from <code>PutRecord</code> and <code>PutRecords</code> operations. Minimum, Maximum, and Average statistics represent the bytes in a single put operation for the stream in the specified time period.</p> <p>Shard-level metric name: <code>IncomingBytes</code></p> <p>Dimensions: <code>StreamName</code></p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Bytes</p>
IncomingRecords	<p>The number of records successfully put to the Amazon Kinesis stream over the specified time period. This metric includes record counts from <code>PutRecord</code> and <code>PutRecords</code> operations. Minimum, Maximum, and Average statistics represent the records in a single put operation for the stream in the specified time period.</p> <p>Shard-level metric name: <code>IncomingRecords</code></p> <p>Dimensions: <code>StreamName</code></p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Count</p>

Metric	Description
PutRecord.Bytes	<p>The number of bytes put to the Amazon Kinesis stream using the PutRecord operation over the specified time period.</p> <p>Dimensions: StreamName</p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Bytes</p>
PutRecord.Latency	<p>The time taken per PutRecord operation, measured over the specified time period.</p> <p>Dimensions: StreamName</p> <p>Statistics: Minimum, Maximum, Average</p> <p>Units: Milliseconds</p>
PutRecord.Success	<p>The number of successful PutRecord operations per Amazon Kinesis stream, measured over the specified time period. Average reflects the percentage of successful writes to a stream.</p> <p>Dimensions: StreamName</p> <p>Statistics: Average, Sum, Samples</p> <p>Units: Count</p>
PutRecords.Bytes	<p>The number of bytes put to the Amazon Kinesis stream using the PutRecords operation over the specified time period.</p> <p>Dimensions: StreamName</p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Bytes</p>
PutRecords.Latency	<p>The time taken per PutRecords operation, measured over the specified time period.</p> <p>Dimensions: StreamName</p> <p>Statistics: Minimum, Maximum, Average</p> <p>Units: Milliseconds</p>
PutRecords.Records	<p>The number of successful records in a PutRecords operation per Amazon Kinesis stream, measured over the specified time period.</p> <p>Dimensions: StreamName</p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Count</p>

Metric	Description
<code>PutRecords.Success</code>	<p>The number of <code>PutRecords</code> operations where at least one record succeeded, per Amazon Kinesis stream, measured over the specified time period.</p> <p>Dimensions: <code>StreamName</code></p> <p>Statistics: Average, Sum, Samples</p> <p>Units: Count</p>
<code>ReadProvisionedThroughputExceeded</code>	<p>The number of <code>GetRecords</code> calls throttled for the stream over the specified time period. The most commonly used statistic for this metric is Average.</p> <p>When the Minimum statistic has a value of 1, all records were throttled for the stream during the specified time period.</p> <p>When the Maximum statistic has a value of 0 (zero), no records were throttled for the stream during the specified time period.</p> <p>Shard-level metric name: <code>ReadProvisionedThroughputExceeded</code></p> <p>Dimensions: <code>StreamName</code></p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Count</p>
<code>WriteProvisionedThroughputExceeded</code>	<p>The number of records rejected due to throttling for the stream over the specified time period. This metric includes throttling from <code>PutRecord</code> and <code>PutRecords</code> operations. The most commonly used statistic for this metric is Average.</p> <p>When the Minimum statistic has a non-zero value, records were being throttled for the stream during the specified time period.</p> <p>When the Maximum statistic has a value of 0 (zero), no records were being throttled for the stream during the specified time period.</p> <p>Shard-level metric name: <code>WriteProvisionedThroughputExceeded</code></p> <p>Dimensions: <code>StreamName</code></p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Count</p>

Enhanced Shard-level Metrics

The `AWS/Kinesis` namespace includes the following shard-level metrics.

Amazon Kinesis sends the following shard-level metrics to CloudWatch every minute. These metrics are not enabled by default. There is a nominal charge for enhanced metrics emitted from Amazon Kinesis. For more information, see [Amazon CloudWatch Pricing](#).

Metric	Description
IncomingBytes	<p>The number of bytes successfully put to the shard over the specified time period. This metric includes bytes from <code>PutRecord</code> and <code>PutRecords</code> operations. Minimum, Maximum, and Average statistics represent the bytes in a single put operation for the shard in the specified time period.</p> <p>Stream-level metric name: <code>IncomingBytes</code></p> <p>Dimensions: <code>StreamName</code>, <code>ShardID</code></p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Bytes</p>
IncomingRecords	<p>The number of records successfully put to the shard over the specified time period. This metric includes record counts from <code>PutRecord</code> and <code>PutRecords</code> operations. Minimum, Maximum, and Average statistics represent the records in a single put operation for the shard in the specified time period.</p> <p>Stream-level metric name: <code>IncomingRecords</code></p> <p>Dimensions: <code>StreamName</code>, <code>ShardID</code></p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Count</p>
IteratorAgeMilliseconds	<p>The age of the last record in all <code>GetRecords</code> calls made against a shard, measured over the specified time period. Age is the difference between the current time and when the last record of the <code>GetRecords</code> call was written to the stream. The Minimum and Maximum statistics can be used to track the progress of Amazon Kinesis consumer applications. A value of 0 (zero) indicates that the records being read are completely caught up with the stream.</p> <p>Stream-level metric name: <code>GetRecords.IteratorAgeMilliseconds</code></p> <p>Dimensions: <code>StreamName</code>, <code>ShardID</code></p> <p>Statistics: Minimum, Maximum, Average, Samples</p> <p>Units: Milliseconds</p>
OutgoingBytes	<p>The number of bytes retrieved from the shard, measured over the specified time period. Minimum, Maximum, and Average statistics represent the bytes in a single <code>GetRecords</code> operation for the shard in the specified time period.</p> <p>Stream-level metric name: <code>GetRecords.Bytes</code></p> <p>Dimensions: <code>StreamName</code>, <code>ShardID</code></p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Bytes</p>

Metric	Description
OutgoingRecords	<p>The number of records retrieved from the shard, measured over the specified time period. Minimum, Maximum, and Average statistics represent the records in a single <code>GetRecords</code> operation for the shard in the specified time period.</p> <p>Stream-level metric name: <code>GetRecords.Records</code></p> <p>Dimensions: <code>StreamName</code>, <code>ShardID</code></p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Count</p>
ReadProvisionedThroughputExceeded	<p>The number of <code>GetRecords</code> calls throttled for the shard over the specified time period. This exception count covers all dimensions of the following limits: 5 reads per shard per second or 2 MB per second per shard. The most commonly used statistic for this metric is Average.</p> <p>When the Minimum statistic has a value of 1, all records were throttled for the shard during the specified time period.</p> <p>When the Maximum statistic has a value of 0 (zero), no records were throttled for the shard during the specified time period.</p> <p>Stream-level metric name: <code>ReadProvisionedThroughputExceeded</code></p> <p>Dimensions: <code>StreamName</code>, <code>ShardID</code></p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Count</p>
WriteProvisionedThroughputExceeded	<p>The number of records rejected due to throttling for the shard over the specified time period. This metric includes throttling from <code>PutRecord</code> and <code>PutRecords</code> operations and covers all dimensions of the following limits: 1,000 records per second per shard or 1 MB per second per shard. The most commonly used statistic for this metric is Average.</p> <p>When the Minimum statistic has a non-zero value, records were being throttled for the shard during the specified time period.</p> <p>When the Maximum statistic has a value of 0 (zero), no records were being throttled for the shard during the specified time period.</p> <p>Stream-level metric name: <code>WriteProvisionedThroughputExceeded</code></p> <p>Dimensions: <code>StreamName</code>, <code>ShardID</code></p> <p>Statistics: Minimum, Maximum, Average, Sum, Samples</p> <p>Units: Count</p>

Dimensions for Amazon Kinesis Metrics

You can use the following dimensions to filter the metrics for Amazon Kinesis Streams.

Dimension	Description
StreamName	The name of the Amazon Kinesis stream.
ShardID	The shard ID within the Amazon Kinesis stream.

AWS Key Management Service Metrics and Dimensions

When you use AWS Key Management Service (AWS KMS) to [import key material](#) into a customer master key (CMK) and set it to expire, AWS KMS sends metrics and dimensions to CloudWatch. For more information, see [Monitoring with Amazon CloudWatch](#) in the *AWS Key Management Service Developer Guide*.

AWS KMS Metrics

The `AWS/KMS` namespace includes the following metrics.

SecondsUntilKeyMaterialExpiration

This metric tracks the number of seconds remaining until imported key material expires. This metric is valid only for CMKs whose origin is `EXTERNAL` and whose key material is or was set to expire. The most useful statistic for this metric is `Minimum`, which tells you the smallest amount of time remaining for all data points in the specified statistic period. The only valid unit for this metric is `Seconds`.

Use this metric to track the amount of time that remains until your imported key material expires. When that amount of time falls below a threshold that you define, you might want to take action such as reimporting the key material with a new expiration date. You can create a CloudWatch alarm to notify you when that happens. For more information, see [Creating CloudWatch Alarms to Monitor AWS KMS Metrics](#) in the *AWS Key Management Service Developer Guide*.

Dimensions for AWS KMS Metrics

AWS KMS metrics use the `AWS/KMS` namespace and have only one valid dimension: `KeyId`. You can use this dimension to view metric data for a specific CMK or set of CMKs.

AWS Lambda Metrics and Dimensions

AWS Lambda sends metrics to CloudWatch every minute. For more information, see [Troubleshooting and Monitoring AWS Lambda Functions with Amazon CloudWatch](#) in the *AWS Lambda Developer Guide*.

AWS Lambda CloudWatch Metrics

The `AWS/Lambda` namespace includes the following metrics.

Metric	Description
Invocations	Measures the number of times a function is invoked in response to an event or invocation API call. This replaces the deprecated <code>RequestCount</code>

Metric	Description
	<p>metric. This includes successful and failed invocations, but does not include throttled attempts. This equals the billed requests for the function. Note that AWS Lambda only sends these metrics to CloudWatch if they have a nonzero value.</p> <p>Units: Count</p>
Errors	<p>Measures the number of invocations that failed due to errors in the function (response code 4XX). This replaces the deprecated <code>ErrorCount</code> metric. Failed invocations may trigger a retry attempt that succeeds. This includes:</p> <ul style="list-style-type: none"> • Handled exceptions (e.g., <code>context.fail(error)</code>) • Unhandled exceptions causing the code to exit • Out of memory exceptions • Timeouts • Permissions errors <p>This does not include invocations that fail due to invocation rates exceeding default concurrent limits (error code 429) or failures due to internal service errors (error code 500).</p> <p>Units: Count</p>
Duration	<p>Measures the elapsed wall clock time from when the function code starts executing as a result of an invocation to when it stops executing. This replaces the deprecated <code>Latency</code> metric. The maximum data point value possible is the function timeout configuration. The billed duration will be rounded up to the nearest 100 millisecond. Note that AWS Lambda only sends these metrics to CloudWatch if they have a nonzero value.</p> <p>Units: Milliseconds</p>
Throttles	<p>Measures the number of Lambda function invocation attempts that were throttled due to invocation rates exceeding the customer's concurrent limits (error code 429). Failed invocations may trigger a retry attempt that succeeds.</p> <p>Units: Count</p>

Errors/Invocations Ratio

When calculating the error rate on Lambda function invocations, it's important to distinguish between an invocation request and an actual invocation. It is possible for the error rate to exceed the number of billed Lambda function invocations. Lambda reports an invocation metric only if the Lambda function code is executed. If the invocation request yields a throttling or other initialization error that prevents the Lambda function code from being invoked, Lambda will report an error, but it does not log an invocation metric.

- Lambda emits `Invocations=1` when the function is executed. If the Lambda function is not executed, nothing is emitted.
- Lambda emits a data point for `Errors` for each invoke request. `Errors=0` means that there is no function execution error. `Errors=1` means that there is a function execution error.
- Lambda emits a data point for `Throttles` for each invoke request. `Throttles=0` means there is no invocation throttle. `Throttles=1` means there is an invocation throttle.

Dimensions for AWS Lambda Metrics

Lambda data can be filtered along any of the following dimensions in the table below.

AWS Lambda CloudWatch Dimensions

You can use the dimensions in the following table to refine the metrics returned for your Lambda functions.

Dimension	Description
FunctionName	Filters the metric data by Lambda function.
Resource	Filters the metric data by Lambda function resource.
Version	Filters the metric data by Lambda version.
Alias	Filters the metric data by Lambda alias.

Amazon Machine Learning Metrics and Dimensions

Amazon Machine Learning sends metrics to CloudWatch every five minutes. For more information, see [Monitoring Amazon ML with Amazon CloudWatch Metrics](#) in the *Amazon Machine Learning Developer Guide*.

Amazon ML Metrics

The AWS/ML namespace includes the following metrics.

Metric	Description
PredictCount	The number of observations received by Amazon ML, measured over the specified time period. Units: Count
PredictFailureCount	The number of invalid or malformed observations received by Amazon ML, measured over the specified time period. Units: Count

Dimensions for Amazon Machine Learning Metrics

Amazon ML data can be filtered along any of the following dimensions in the table below.

Dimension	Description
MLModelId	The identifier of an Amazon ML model. All available statistics are filtered by <code>MLModelId</code> .
RequestMode	An indicator specifying whether observations were received as part of a batch prediction request or as real-time predict requests. All available statistics are filtered by <code>RequestMode</code> .

AWS OpsWorks Metrics and Dimensions

AWS OpsWorks sends metrics to CloudWatch for each active stack every minute. Detailed monitoring is enabled by default. For more information, see [Monitoring](#) in the *AWS OpsWorks User Guide*.

AWS OpsWorks Metrics

The `AWS/OpsWorks` namespace includes the following metrics.

Metric	Description
<code>cpu_idle</code>	The percentage of time that the CPU is idle. Units: Percent
<code>cpu_nice</code>	The percentage of time that the CPU is handling processes with a positive nice value, which have lower scheduling priority. For information, see nice (Unix) . Units: Percent
<code>cpu_system</code>	The percentage of time that the CPU is handling system operations. Units: Percent
<code>cpu_user</code>	The percentage of time that the CPU is handling user operations. Units: Percent
<code>cpu_waitio</code>	The percentage of time that the CPU is waiting for input/output operations. Units: Percent
<code>load_1</code>	The load averaged over a 1-minute window. Units: Unix load units
<code>load_5</code>	The load averaged over a 5-minute window. Units: Unix load units
<code>load_15</code>	The load averaged over a 15-minute window. Units: Unix load units
<code>memory_buffers</code>	The amount of buffered memory. Units: Kilobytes
<code>memory_cached</code>	The amount of cached memory. Units: Kilobytes
<code>memory_free</code>	The amount of free memory. Units: Kilobytes
<code>memory_swap</code>	The amount of swap space. Units: Kilobytes

Metric	Description
memory_total	The total amount of memory. Units: Kilobytes
memory_used	The amount of memory in use. Units: Kilobytes
procs	The number of active processes. Units: Count

Dimensions for AWS OpsWorks Metrics

AWS OpsWorks data can be filtered along any of the following dimensions in the table below.

Dimension	Description
StackId	Average values for a stack.
LayerId	Average values for a layer.
InstanceId	Average values for an instance.

Amazon Polly Metrics

Amazon Polly sends metrics to CloudWatch. For more information, see the *Amazon Polly Developer Guide*.

Amazon Polly Metrics

Amazon Polly produces the following metrics for each request. These metrics are aggregated and in one minute intervals sent to CloudWatch where they are available in the `AWS/Polly` namespace.

Metric	Description
RequestCharacters	The number of characters in the request. This is billable characters only and does not include SSML tags. Valid Dimension: Operation Valid Statistics: Minimum, Maximum, Average, SampleCount, Sum Unit: Count
ResponseLatency	The latency between when the request was made and the start of the streaming response. Valid Dimensions: Operation Valid Statistics: Minimum, Maximum, Average, SampleCount Unit: milliseconds

Metric	Description
2XXCount	<p>HTTP 200 level code returned upon a successful response.</p> <p>Valid Dimensions: Operation</p> <p>Valid Statistics: Average, SampleCount, Sum</p> <p>Unit: Count</p>
4XXCount	<p>HTTP 400 level error code returned upon an error. For each successful response, a zero (0) is emitted.</p> <p>Valid Dimensions: Operation</p> <p>Valid Statistics: Average, SampleCount, Sum</p> <p>Unit: Count</p>
5XXCount	<p>HTTP 500 level error code returned upon an error. For each successful response, a zero (0) is emitted.</p> <p>Valid Dimensions: Operation</p> <p>Valid Statistics: Average, SampleCount, Sum</p> <p>Unit: Count</p>

Dimensions for Amazon Polly Metrics

Amazon Polly provides metrics for the following dimension.

Dimension	Description
Operation	<p>Metrics are grouped by the API method they refer to. Possible values are <code>SynthesizeSpeech</code>, <code>PutLexicon</code>, <code>DescribeVoices</code>, etc.</p>

Amazon Redshift Metrics and Dimensions

Amazon Redshift sends metrics to CloudWatch for each active cluster every minute. Detailed monitoring is enabled by default. For more information, see [Monitoring Amazon Redshift Cluster Performance](#) in the *Amazon Redshift Cluster Management Guide*.

Amazon Redshift Metrics

The `AWS/Redshift` namespace includes the following metrics.

Metric	Description
CPUUtilization	<p>The percentage of CPU utilization. For clusters, this metric represents an aggregation of all nodes (leader and compute) CPU utilization values.</p> <p>Units: Percent</p>

Metric	Description
	Dimensions: NodeID, ClusterIdentifier
DatabaseConnections	<p>The number of database connections to a cluster.</p> <p>Units: Count</p> <p>Dimensions: ClusterIdentifier</p>
HealthStatus	<p>Indicates the health of the cluster. Every minute the cluster connects to its database and performs a simple query. If it is able to perform this operation successfully, the cluster is considered healthy. Otherwise, the cluster is unhealthy. An unhealthy status can occur when the cluster database is under extremely heavy load or if there is a configuration problem with a database on the cluster. The exception to this is when the cluster is undergoing maintenance. Even though your cluster might be unavailable due to maintenance tasks, the cluster remains in HEALTHY state. For more information, see Maintenance Windows in the <i>Amazon Redshift Cluster Management Guide</i>.</p> <p>Note In Amazon CloudWatch this metric is reported as 1 or 0 whereas in the Amazon CloudWatch console, this metric is displayed with the words HEALTHY or UNHEALTHY for convenience. When this metric is displayed in the Amazon CloudWatch console, sampling averages are ignored and only HEALTHY or UNHEALTHY are displayed. In Amazon CloudWatch, values different than 1 and 0 may occur because of sampling issue. Any value below 1 for HealthStatus is reported as 0 (UNHEALTHY).</p> <p>Units: 1/0 (HEALTHY/UNHEALTHY in the Amazon CloudWatch console)</p> <p>Dimensions: ClusterIdentifier</p>
MaintenanceMode	<p>Indicates whether the cluster is in maintenance mode.</p> <p>Note In Amazon CloudWatch this metric is reported as 1 or 0 whereas in the Amazon CloudWatch console, this metric is displayed with the words ON or OFF for convenience. When this metric is displayed in the Amazon CloudWatch console, sampling averages are ignored and only ON or OFF are displayed. In Amazon CloudWatch, values different than 1 and 0 may occur because of sampling issues. Any value greater than 0 for MaintenanceMode is reported as 1 (ON).</p> <p>Units: 1/0 (ON/OFF in the Amazon CloudWatch console).</p> <p>Dimensions: ClusterIdentifier</p>
NetworkReceiveThroughput	<p>The rate at which the node or cluster receives data.</p> <p>Units: Bytes/seconds (MB/s in the Amazon CloudWatch console)</p> <p>Dimensions: NodeID, ClusterIdentifier</p>

Metric	Description
NetworkTransmitThroughput	The rate at which the node or cluster writes data. Units: Bytes/second (MB/s in the Amazon CloudWatch console) Dimensions: NodeID, ClusterIdentifier
PercentageDiskSpaceUsed	The percent of disk space used. Units: Percent Dimensions: NodeID, ClusterIdentifier
ReadIOPS	The average number of disk read operations per second. Units: Count/second Dimensions: NodeID
ReadLatency	The average amount of time taken for disk read I/O operations. Units: Seconds Dimensions: NodeID
ReadThroughput	The average number of bytes read from disk per second. Units: Bytes (GB/s in the Amazon CloudWatch console) Dimensions: NodeID
WriteIOPS	The average number of write operations per second. Units: Count/seconds Dimensions: NodeID
WriteLatency	The average amount of time taken for disk write I/O operations. Units: Seconds Dimensions: NodeID
WriteThroughput	The average number of bytes written to disk per second. Units: Bytes (GB/s in the Amazon CloudWatch console) Dimensions: NodeID

Dimensions for Amazon Redshift Metrics

Amazon Redshift data can be filtered along any of the following dimensions in the table below.

Dimension	Description
NodeID	Filters requested data that is specific to the nodes of a cluster. NodeID will be either "Leader", "Shared", or "Compute-N" where N is 0, 1, ... for the number of nodes in the cluster. "Shared"

Dimension	Description
	means that the cluster has only one node, i.e. the leader node and compute node are combined.
ClusterIdentifier	Filters requested data that is specific to the cluster. Metrics that are specific to clusters include <code>HealthStatus</code> , <code>MaintenanceMode</code> , and <code>DatabaseConnections</code> . In general metrics in for this dimension (e.g. <code>ReadIOPS</code>) that are also metrics of nodes represent an aggregate of the node metric data. You should take care in interpreting these metrics because they aggregate behavior of leader and compute nodes.

Amazon RDS Metrics and Dimensions

Amazon Relational Database Service sends metrics to CloudWatch for each active database instance every minute. Detailed monitoring is enabled by default. For more information, see [Monitoring a DB Instance](#) in the *Amazon Relational Database Service User Guide*.

Amazon RDS Metrics

The `AWS/RDS` namespace includes the following metrics.

Metric	Description
<code>BinLogDiskUsage</code>	The amount of disk space occupied by binary logs on the master. Applies to MySQL read replicas. Units: Bytes
<code>CPUUtilization</code>	The percentage of CPU utilization. Units: Percent
<code>CPUCreditUsage</code>	[T2 instances] The number of CPU credits consumed during the specified period. This metric identifies the amount of time during which physical CPUs were used for processing instructions by virtual CPUs allocated to the instance. CPU Credit metrics are available at a 5 minute frequency. Units: Count
<code>CPUCreditBalance</code>	[T2 instances] The number of CPU credits that an instance has accumulated. This metric determines how long an instance can burst beyond its baseline performance level at a given rate. CPU Credit metrics are available at a 5 minute frequency. Units: Count
<code>DatabaseConnections</code>	The number of database connections in use. Units: Count

Metric	Description
DiskQueueDepth	The number of outstanding IOs (read/write requests) waiting to access the disk. Units: Count
FreeableMemory	The amount of available random access memory. Units: Bytes
FreeStorageSpace	The amount of available storage space. Units: Bytes
ReplicaLag	The amount of time a Read Replica DB instance lags behind the source DB instance. Applies to MySQL, MariaDB, and PostgreSQL Read Replicas. Units: Seconds
SwapUsage	The amount of swap space used on the DB instance. Units: Bytes
ReadIOPS	The average number of disk I/O operations per second. Units: Count/Second
WriteIOPS	The average number of disk I/O operations per second. Units: Count/Second
ReadLatency	The average amount of time taken per disk I/O operation. Units: Seconds
WriteLatency	The average amount of time taken per disk I/O operation. Units: Seconds
ReadThroughput	The average number of bytes read from disk per second. Units: Bytes/Second
WriteThroughput	The average number of bytes written to disk per second. Units: Bytes/Second
NetworkReceiveThroughput	The incoming (Receive) network traffic on the DB instance, including both customer database traffic and Amazon RDS traffic used for monitoring and replication. Units: Bytes/second
NetworkTransmitThroughput	The outgoing (Transmit) network traffic on the DB instance, including both customer database traffic and Amazon RDS traffic used for monitoring and replication. Units: Bytes/second

Dimensions for RDS Metrics

Amazon RDS data can be filtered along any of the following dimensions in the table below.

Dimension	Description
<code>DBInstanceIdentifier</code>	This dimension filters the data you request for a specific database instance.
<code>DBClusterIdentifier</code>	This dimension filters the data you request for a specific Amazon Aurora DB cluster.
<code>DatabaseClass</code>	This dimension filters the data you request for all instances in a database class. For example, you can aggregate metrics for all instances that belong to the database class <code>db.m1.small</code>
<code>EngineName</code>	This dimension filters the data you request for the identified engine name only. For example, you can aggregate metrics for all instances that have the engine name <code>mysql</code> .

Amazon Route 53 Metrics and Dimensions

Amazon Route 53 sends metrics to CloudWatch. CloudWatch provides detailed monitoring of Amazon Route 53 by default. Amazon Route 53 sends one-minute metrics to CloudWatch. For more information, see [Monitoring Health Checks Using Amazon CloudWatch](#) in the *Amazon Route 53 Developer Guide*.

Note

To get Amazon Route 53 metrics using CloudWatch, you must choose US East (N. Virginia) as the region. Amazon Route 53 metrics are not available if you select any other region. You can also optionally specify a `Region` dimension. For more information, see [Dimensions for Amazon Route 53 Metrics](#) (p. 125).

Amazon Route 53 Metrics

The `AWS/Route53` namespace includes the following metrics.

Metric	Description
<code>ChildHealthCheckHealthyCount</code>	<p>For a calculated health check, the number of health checks that are healthy among the health checks that Amazon Route 53 is monitoring.</p> <p>Valid statistics: Average (recommended), Minimum, Maximum</p> <p>Units: Healthy health checks</p>
<code>ConnectionTime</code>	<p>The average time, in milliseconds, that it took Amazon Route 53 health checkers to establish a TCP connection with the endpoint. You can view <code>ConnectionTime</code> for a health check either across all regions or for a selected geographic region.</p> <p>Valid statistics: Average (recommended), Minimum, Maximum</p> <p>Units: Milliseconds</p>

Metric	Description
HealthCheckPercentageHealthy	<p>The percentage of Amazon Route 53 health checkers that consider the selected endpoint to be healthy. You can view <code>HealthCheckPercentageHealthy</code> only across all regions; data is not available for a selected region.</p> <p>Valid statistics: Average, Minimum, Maximum</p> <p>Units: Percent</p>
HealthCheckStatus	<p>The status of the health check endpoint that CloudWatch is checking. 1 indicates healthy, and 0 indicates unhealthy. You can view <code>HealthCheckStatus</code> only across all regions; data is not available for a selected region.</p> <p>Valid statistics: Minimum</p> <p>Units: none</p>
SSLHandshakeTime	<p>The average time, in milliseconds, that it took Amazon Route 53 health checkers to complete the SSL handshake. You can view <code>SSLHandshakeTime</code> for a health check either across all regions or for a selected geographic region.</p> <p>Valid statistics: Average (recommended), Minimum, Maximum</p> <p>Units: Milliseconds</p>
TimeToFirstByte	<p>The average time, in milliseconds, that it took Amazon Route 53 health checkers to receive the first byte of the response to an HTTP or HTTPS request. You can view <code>TimeToFirstByte</code> for a health check either across all regions or for a selected geographic region.</p> <p>Valid statistics: Average (recommended), Minimum, Maximum</p> <p>Units: Milliseconds</p>

Dimensions for Amazon Route 53 Metrics

Amazon Route 53 metrics use the `AWS/Route53` namespace and provide metrics for `HealthCheckId`. When retrieving metrics, you must supply the `HealthCheckId` dimension.

In addition, for `ConnectionTime`, `SSLHandshakeTime`, and `TimeToFirstByte`, you can optionally specify `Region`. If you omit `Region`, CloudWatch returns metrics across all regions. If you include `Region`, CloudWatch returns metrics only for the specified region.

For more information, see [Monitoring Health Checks Using CloudWatch](#) in the *Amazon Route 53 Developer Guide*.

Amazon Simple Email Service Metrics and Dimensions

Amazon Simple Email Service sends data points to CloudWatch for email sending events. For more information, see [Retrieving Amazon SES Event Data from CloudWatch](#) in the *Amazon Simple Email Service Developer Guide*.

Amazon SES Event Metrics

The following metrics are available from Amazon SES.

Metric	Description
Bounce	The recipient's mail server permanently rejected the email. This event corresponds to hard bounces. Soft bounces are included only when Amazon SES fails to deliver the email after retrying for a period of time. Unit: count
Complaint	The recipient marked the email as spam. Unit: count
Delivery	Amazon SES successfully delivered the email to the recipient's mail server. Unit: count
Reject	Amazon SES initially accepted the email, but later rejected it because the email contained a virus. Unit: count
Send	The email sending API call to Amazon SES was successful and Amazon SES will attempt to deliver the email. Unit: count

Dimensions for Amazon SES Metrics

CloudWatch uses the dimension names that you specify when you add a CloudWatch event destination to a configuration set in Amazon SES. For more information, see [Set Up a CloudWatch Event Destination for Amazon SES Event Publishing](#).

Amazon Simple Notification Service Metrics and Dimensions

Amazon Simple Notification Service sends data points to CloudWatch for several metrics. All active topics automatically send five-minute metrics to CloudWatch. Detailed monitoring, or one-minute

metrics, is currently unavailable for Amazon Simple Notification Service. A topic stays active for six hours from the last activity (for example, any API call) on the topic. For more information, see [Monitoring Amazon SNS with Amazon CloudWatch](#) in the *Amazon Simple Notification Service Developer Guide*.

Amazon Simple Notification Service Metrics

The AWS/SNS namespace includes the following metrics.

Metric	Description
NumberOfMessagesPublished	The number of messages published. Units: Count Valid Statistics: Sum
PublishSize	The size of messages published. Units: Bytes Valid Statistics: Minimum, Maximum, Average and Count
NumberOfNotificationsDelivered	The number of messages successfully delivered. Units: Count Valid Statistics: Sum
NumberOfNotificationsFailed	The number of messages that Amazon SNS failed to deliver. This metric is applied after Amazon SNS stops attempting message deliveries to Amazon SQS, email, SMS, or mobile push endpoints. Each delivery attempt to an HTTP or HTTPS endpoint adds 1 to the metric. For all other endpoints, the count increases by 1 when the message is not delivered (regardless of the number of attempts). You can control the number of retries for HTTP endpoints; for more information, see Setting Amazon SNS Delivery Retry Policies for HTTP/HTTPS Endpoints . Units: Count Valid Statistics: Sum, Average
SMSSuccessRate	The rate of successful SMS message deliveries. Units: Count Valid Statistics: Sum, Average, Data Samples

Dimensions for Amazon Simple Notification Service Metrics

Amazon SNS sends the following dimensions to CloudWatch.

Dimension	Description
Application	Filters on application objects, which represent an app and device registered with one of the supported push notification services, such as APNS and GCM.
Application,Platform	Filters on application and platform objects, where the platform objects are for the supported push notification services, such as APNS and GCM.
Country	Filters on the destination country of an SMS message. The country is represented by its ISO 3166-1 alpha-2 code.
Platform	Filters on platform objects for the push notification services, such as APNS and GCM.
TopicName	Filters on Amazon SNS topic names.
SMSType	Filters on the message type of SMS message. Can be <i>promotional</i> or <i>transactional</i> .

Amazon SQS Metrics and Dimensions

Amazon SQS sends data points to CloudWatch for several metrics. All active queues automatically send five-minute metrics to CloudWatch. Detailed monitoring, or one-minute metrics, is currently unavailable for Amazon SQS. A queue stays active for six hours from the last activity (for example, any API call) on the queue. For more information, see [Monitoring Amazon SQS with Amazon CloudWatch](#) in the *Amazon Simple Queue Service Developer Guide*.

Amazon SQS Metrics

The AWS/SQS namespace includes the following metrics.

Metric	Description
ApproximateAgeOfOldestMessage	<p>The approximate age of the oldest non-deleted message in the queue.</p> <p>Units: <i>Seconds</i></p> <p>Valid Statistics: Average, Minimum, Maximum, Sum, Data Samples (displays as Sample Count in the Amazon SQS console)</p>
ApproximateNumberOfMessagesDelayed	<p>The number of messages in the queue that are delayed and not available for reading immediately. This can happen when the queue is configured as a delay queue or when a message has been sent with a delay parameter.</p> <p>Units: <i>Count</i></p> <p>Valid Statistics: Average, Minimum, Maximum, Sum, Data Samples (displays as Sample Count in the Amazon SQS console)</p>

Metric	Description
ApproximateNumberOfMessagesNotVisible	<p>The number of messages that are "in flight." Messages are considered in flight if they have been sent to a client but have not yet been deleted or have not yet reached the end of their visibility window.</p> <p>Units: <i>Count</i></p> <p>Valid Statistics: Average, Minimum, Maximum, Sum, Data Samples (displays as Sample Count in the Amazon SQS console)</p>
ApproximateNumberOfMessagesVisible	<p>The number of messages available for retrieval from the queue.</p> <p>Units: <i>Count</i></p> <p>Valid Statistics: Average, Minimum, Maximum, Sum, Data Samples (displays as Sample Count in the Amazon SQS console)</p>
NumberOfEmptyReceives	<p>The number of <code>ReceiveMessage</code> API calls that did not return a message.</p> <p>Units: <i>Count</i></p> <p>Valid Statistics: Average, Minimum, Maximum, Sum, Data Samples (displays as Sample Count in the Amazon SQS console)</p>
NumberOfMessagesDeleted	<p>The number of messages deleted from the queue.</p> <p>Units: <i>Count</i></p> <p>Valid Statistics: Average, Minimum, Maximum, Sum, Data Samples (displays as Sample Count in the Amazon SQS console)</p>
NumberOfMessagesReceived	<p>The number of messages returned by calls to the <code>ReceiveMessage</code> API action.</p> <p>Units: <i>Count</i></p> <p>Valid Statistics: Average, Minimum, Maximum, Sum, Data Samples (displays as Sample Count in the Amazon SQS console)</p>
NumberOfMessagesSent	<p>The number of messages added to a queue.</p> <p>Units: <i>Count</i></p> <p>Valid Statistics: Average, Minimum, Maximum, Sum, Data Samples (displays as Sample Count in the Amazon SQS console)</p>

Metric	Description
SentMessageSize	<p>The size of messages added to a queue.</p> <p>Units: <i>Bytes</i></p> <p>Valid Statistics: Average, Minimum, Maximum, Sum, Data Samples (displays as Sample Count in the Amazon SQS console)</p> <p>Note that <code>SentMessageSize</code> does not display as an available metric in the CloudWatch console until at least one message is sent to the corresponding queue.</p>

Dimensions for Amazon SQS Metrics

The only dimension that Amazon SQS sends to CloudWatch is `QueueName`. This means that all available statistics are filtered by `QueueName`.

Amazon Simple Storage Service Metrics and Dimensions

Amazon Simple Storage Service sends data points to CloudWatch for several metrics, such as object counts and bytes stored, once a day. For more information, see [Monitoring Amazon S3 with CloudWatch](#) in the *Amazon Simple Storage Service Developer Guide*.

Amazon S3 CloudWatch Metrics

The `AWS/S3` namespace includes the following daily storage metrics for buckets.

Metric	Description
BucketSizeBytes	<p>The amount of data in bytes stored in a bucket in the Standard storage class, Standard - Infrequent Access (Standard_IA) storage class, or the Reduced Redundancy Storage (RRS) storage class.</p> <p>Valid storage type filters: <code>StandardStorage</code>, or <code>StandardIAStorage</code>, or <code>ReducedRedundancyStorage</code> (see <code>StorageType</code> dimension)</p> <p>Units: Bytes</p> <p>Valid statistics: Average</p>
NumberOfObjects	<p>The total number of objects stored in a bucket for all storage classes except for the GLACIER storage class.</p> <p>Valid storage type filters: <code>AllStorageTypes</code> only (see <code>StorageType</code> dimension)</p> <p>Units: Count</p> <p>Valid statistics: Average</p>

The `AWS/S3` namespace includes the following request metrics.

Metric	Description
<code>AllRequests</code>	<p>The total number of HTTP requests made to a bucket, regardless of type. If you use a metrics configuration with a filter, this metric returns only the HTTP requests made to the objects in the bucket that meet the filter requirements.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>
<code>GetRequests</code>	<p>The number of HTTP GET requests made for objects in a bucket. This doesn't include list operations.</p> <p>Paginated list-oriented requests, such as List Multipart Uploads, List Parts, Get Bucket Object Versions, and others, are not included in this metric.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>
<code>PutRequests</code>	<p>The number of HTTP PUT requests made for objects in a bucket.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>
<code>DeleteRequests</code>	<p>The number of HTTP DELETE requests made for objects in a bucket. This also includes Delete Multiple Objects requests.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>
<code>HeadRequests</code>	<p>The number of HTTP HEAD requests made to a bucket.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>
<code>PostRequests</code>	<p>The number of HTTP POST requests made to a bucket.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>
<code>ListRequests</code>	<p>The number of HTTP requests that list the contents of a bucket.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>
<code>BytesDownloaded</code>	<p>The number bytes downloaded for requests made to a bucket, where the response includes a body.</p> <p>Units: Bytes</p> <p>Valid statistics: Average (bytes per request), Sum (bytes per period), Sample Count, Min, Max</p>
<code>BytesUploaded</code>	<p>The number bytes uploaded to a bucket that contain a request body.</p>

Metric	Description
	<p>Units: Bytes</p> <p>Valid statistics: Average (bytes per request), Sum (bytes per period), Sample Count, Min, Max</p>
4xxErrors	<p>The number of HTTP 4xx client error status code requests made to a bucket with a value of 0 or 1. The <code>average</code> statistic shows the error rate, and the <code>sum</code> statistic shows the count of that type of error, during each period.</p> <p>Units: Count</p> <p>Valid statistics: Average (reports per request), Sum (reports per period), Min, Max, Sample Count</p>
5xxErrors	<p>The number of HTTP 5xx server error status code requests made to a bucket with a value of either 0 or 1. The <code>average</code> statistic shows the error rate, and the <code>sum</code> statistic shows the count of that type of error, during each period.</p> <p>Units: Counts</p> <p>Valid statistics: Average (reports per request), Sum (reports per period), Min, Max, Sample Count</p>
FirstByteLatency	<p>The per-request time from the complete request being received by a bucket to when the response starts to be returned.</p> <p>Units: Milliseconds</p> <p>Valid statistics: Average, Sum, Min, Max, Sample Count</p>
TotalRequestLatency	<p>The elapsed per-request time from the first byte received to the last byte sent to a bucket. This includes the time taken to receive the request body and send the response body, which is not included in <code>FirstByteLatency</code>.</p> <p>Units: Milliseconds</p> <p>Valid statistics: Average, Sum, Min, Max, Sample Count</p>

Amazon S3 CloudWatch Dimensions

The following dimensions are used to filter Amazon S3 metrics.

Dimension	Description
BucketName	Filters the data you request for the identified bucket only.
StorageType	Filters the data stored in a bucket by the type of storage. The types are <code>StandardStorage</code> for the Standard storage class, <code>StandardIAStorage</code> for the Standard_IA storage class, <code>ReducedRedundancyStorage</code> for the Reduced Redundancy Storage (RRS) class, and <code>AllStorageTypes</code> . Note that the <code>AllStorageTypes</code> type does not include the <code>GLACIER</code> storage class.

Dimension	Description
FilterId	Filters metrics configurations that you specify for request metrics on a bucket, for example, a prefix or a tag. You specify a filter ID when you create a metrics configuration.

Amazon SWF Metrics and Dimensions

Amazon SWF sends data points to CloudWatch for several metrics. Some of the Amazon SWF metrics for CloudWatch are time intervals, always measured in milliseconds. These metrics generally correspond to stages of your workflow execution for which you can set workflow and activity timeouts, and have similar names. For example, the **DecisionTaskStartToCloseTime** metric measures the time it took for the decision task to complete after it began executing, which is the same time period for which you can set a **DecisionTaskStartToCloseTimeout** value.

Other Amazon SWF metrics report results as a count. For example, **WorkflowsCanceled**, records a result as either one or zero, indicating whether or not the workflow was canceled. A value of zero does not indicate that the metric was not reported, only that the condition described by the metric did not occur. For count metrics, minimum and maximum will always be either zero or one, but average will be a value ranging from zero to one. For more information, see [Viewing Amazon SWF Metrics for CloudWatch using the AWS Management Console](#); in the *Amazon Simple Workflow Service Developer Guide*.

Workflow Metrics

The `AWS/SWF` namespace includes the following metrics for Amazon SWF workflows:

Metric	Description
DecisionTaskScheduleToStartTime	The time interval, in milliseconds, between the time that the decision task was scheduled and the time it was picked up by a worker and started.
DecisionTaskStartToCloseTime	The time interval, in milliseconds, between the time that the decision task was started and the time it was closed.
DecisionTasksCompleted	The count of decision tasks that have been completed.
StartedDecisionTasksTimedOutOnClose	The count of decision tasks that started but timed out on closing.
WorkflowStartToCloseTime	The time, in milliseconds, between the time the workflow started and the time it closed.
WorkflowsCanceled	The count of workflows that were canceled.
WorkflowsCompleted	The count of workflows that completed.
WorkflowsContinuedAsNew	The count of workflows that continued as new.
WorkflowsFailed	the count of workflows that failed.
WorkflowsTerminated	the count of workflows that were terminated.
WorkflowsTimedOut	The count of workflows that timed out, for any reason.

Dimensions for Amazon SWF Workflow Metrics

Dimension	Description
Domain	The Amazon SWF domain that the workflow is running in.
WorkflowTypeName	The name of the workflow type for this workflow execution.
WorkflowTypeVersion	The version of the workflow type for this workflow execution.

Activity Metrics

The `AWS/SWF` namespace includes the following metrics for Amazon SWF activities:

Metric	Description
ActivityTaskScheduleToCloseTime	The time interval, in milliseconds, between the time when the activity was scheduled to when it closed.
ActivityTaskScheduleToStartTime	The time interval, in milliseconds, between the time when the activity task was scheduled and when it started.
ActivityTaskStartToCloseTime	The time interval, in milliseconds, between the time when the activity task started and when it was closed.
ActivityTasksCanceled	The count of activity tasks that were canceled.
ActivityTasksCompleted	The count of activity tasks that completed.
ActivityTasksFailed	The count of activity tasks that failed.
ScheduledActivityTasksTimedOutOnClose	The count of activity tasks that were scheduled but timed out on close.
ScheduledActivityTasksTimedOutOnStart	The count of activity tasks that were scheduled but timed out on start.
StartedActivityTasksTimedOutOnClose	The count of activity tasks that were started but timed out on close.
StartedActivityTasksTimedOutOnHeartbeat	The count of activity tasks that were started but timed out due to a heartbeat timeout.

Dimensions for Amazon SWF Activity Metrics

Dimension	Description
Domain	The Amazon SWF domain that the activity is running in.
ActivityTypeName	The name of the activity type.
ActivityTypeVersion	The version of the activity type

AWS Storage Gateway Metrics and Dimensions

AWS Storage Gateway sends data points to CloudWatch for several metrics. All active queues automatically send five-minute metrics to CloudWatch. Detailed monitoring, or one-minute metrics,

is currently unavailable for AWS Storage Gateway. For more information, see [Monitoring Your AWS Storage Gateway](#) in the *AWS Storage Gateway User Guide*.

AWS Storage Gateway Metrics

The `AWS/StorageGateway` namespace includes the following metrics.

You can use these metrics to get information about your gateways. Specify the `GatewayId` or `GatewayName` dimension for each metric to view the data for a gateway. Note that these metrics are measured in 5-minute intervals.

Metric	Description	Gateway-Cached	Gateway-Stored	Gateway-VTL
CacheHitPercent	Percent of application reads served from the cache. This metric applies only to the gateway-cached volume setup. The sample is taken at the end of the reporting period. Units: Percent	yes	no	yes
CacheUsePercent	Percent use of the gateway's cache storage. This metric applies only to the gateway-cached volume setup. The sample is taken at the end of the reporting period. Units: Percent	yes	no	yes
CacheEvictionPercent	Percent of the gateway's cache that has not been persisted to AWS. This metric applies only to the gateway-cached volume setup. The sample is taken at the end of the reporting period. Units: Percent	yes	no	yes
CloudBytesCompressed	The total number of compressed bytes that the gateway downloaded from AWS during the reporting period.	yes	yes	yes

Metric	Description	Gateway-Cached	Gateway-Stored	Gateway-VTL
	<p>Use this metric with the <code>Sum</code> statistic to measure throughput and with the <code>Samples</code> statistic to measure input/output operations per second (IOPS).</p> <p>Units: Bytes</p>			
CloudDownloadLatency	<p>The total number of milliseconds spent reading data from AWS during the reporting period.</p> <p>Use this metric with the <code>Average</code> statistic to measure latency.</p> <p>Units: Milliseconds</p>	yes	yes	yes
CloudBytesUploaded	<p>The total number of compressed bytes that the gateway uploaded to AWS during the reporting period.</p> <p>Use this metric with the <code>Sum</code> statistic to measure throughput and with the <code>Samples</code> statistic to measure IOPS.</p> <p>Units: Bytes</p>	yes	yes	yes
UploadBufferFree	<p>The total amount of unused space in the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Units: Bytes</p>	yes	no	yes
CacheFree	<p>The total amount of unused space in the gateway's cache storage. The sample is taken at the end of the reporting period.</p> <p>Units: Bytes</p>	yes	no	yes

Metric	Description	Gateway-Cached	Gateway-Stored	Gateway-VTL
UploadBufferUse	<p>The percentage of the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Units: Percent</p>	yes	no	yes
UploadBytes	<p>The total number of bytes being used in the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Units: Bytes</p>	yes	no	yes
CacheUsedBytes	<p>The total number of bytes being used in the gateway's cache storage. The sample is taken at the end of the reporting period.</p> <p>Units: Bytes</p>	yes	no	yes
QueuedBytes	<p>The number of bytes waiting to be written to AWS, sampled at the end of the reporting period for all volumes in the gateway. These bytes are kept in your gateway's working storage.</p> <p>Units: Bytes</p>	yes	yes	yes
ReadBytes	<p>The total number of bytes read from your on-premises applications in the reporting period for all volumes in the gateway.</p> <p>Use this metric with the <code>Sum</code> statistic to measure throughput and with the <code>Samples</code> statistic to measure IOPS.</p> <p>Units: Bytes</p>	yes	yes	yes

Metric	Description	Gateway-Cached	Gateway-Stored	Gateway-VTL
ReadTime	<p>The total number of milliseconds spent to do read operations from your on-premises applications in the reporting period for all volumes in the gateway.</p> <p>Use this metric with the <code>Average</code> statistic to measure latency.</p> <p>Units: Milliseconds</p>	yes	yes	yes
TotalCacheSize	<p>The total size of the cache in bytes. This metric applies only to the gateway-cached volume setup. The sample is taken at the end of the reporting period.</p> <p>Units: Bytes</p>	yes	no	yes
WriteBytes	<p>The total number of bytes written to your on-premises applications in the reporting period for all volumes in the gateway.</p> <p>Use this metric with the <code>Sum</code> statistic to measure throughput and with the <code>Samples</code> statistic to measure IOPS.</p> <p>Units: Bytes</p>	yes	yes	yes

Metric	Description	Gateway-Cached	Gateway-Stored	Gateway-VTL
WriteTime	<p>The total number of milliseconds spent to do write operations from your on-premises applications in the reporting period for all volumes in the gateway.</p> <p>Use this metric with the <i>Average</i> statistic to measure latency.</p> <p>Units: Milliseconds</p>	yes	yes	yes
TimeSinceRecoveryPoint	<p>The time since the last available recovery point.</p> <p>Units: Seconds</p>	yes	yes	no

Metric	Description	Gateway-Cached	Gateway-Stored	Gateway-VTL
WorkingStorageFree	<p>The total amount of unused space in the gateway's working storage. The sample is taken at the end of the reporting period.</p> <p>Note Working storage applies only to the gateway-stored volume setup. The upload buffer applies to both the gateway-stored and gateway-cached volume setups. If you are working with both types of gateway setups, you might find it more convenient to use just the corresponding upload buffer metric, UploadBufferFree.</p> <p>Units: Bytes</p>	no	yes	no

Metric	Description	Gateway-Cached	Gateway-Stored	Gateway-VTL
WorkingStoragePercentUsed	<p>Percent use of the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Note Working storage applies only to the gateway-stored volume setup. The upload buffer applies to both the gateway-stored and gateway-cached volume setups. If you are working with both types of gateway setups, you might find it more convenient to use just the corresponding upload buffer metric, UploadBufferPercentUsed.</p> <p>Units: Percent</p>	no	yes	no

Metric	Description	Gateway-Cached	Gateway-Stored	Gateway-VTL
WorkingStorageUsed	<p>The total number of bytes being used in the gateway's upload buffer. The sample is taken at the end of the reporting period.</p> <p>Note Working storage applies only to the gateway-stored volume setup. The upload buffer applies to both the gateway-stored and gateway-cached volume setups. If you are working with both types of gateway setups, you might find it more convenient to use just the corresponding upload buffer metric, UploadBufferUsed.</p> <p>Units: Bytes</p>	no	yes	no

The following table describes the AWS Storage Gateway metrics that you can use to get information about your storage volumes. Specify the `VolumeId` dimension for each metric to view the data for a storage volume.

Metric	Description	Gateway-Cached	Gateway-Stored
CacheHitPercent	<p>Percent of application read operations from the volume that are served from cache. This metric applies only to cached volumes. The sample is taken at the end of the reporting period.</p> <p>When there are no application read operations from the volume, this metric reports 100 percent.</p> <p>Units: Percent</p>	yes	no
CachePercentUsed	<p>The volume's contribution to the overall percent use of the gateway's cache storage. This metric applies only to cached volumes. The sample is taken at the end of the reporting period.</p> <p>Use the <code>CachePercentUsed</code> metric of the gateway to view overall percent use of the gateway's cache storage.</p> <p>Units: Percent</p>	yes	no
CachePercentDirty	<p>The volume's contribution to the overall percentage of the gateway's cache that has not been persisted to AWS. This metric applies only to volumes in a gateway-cached setup. The sample is taken at the end of the reporting period.</p> <p>Use the <code>CachePercentDirty</code> metric of the gateway to view the overall percentage of the gateway's cache that has not been persisted to AWS.</p> <p>Units: Percent</p>	yes	no

Metric	Description	Gateway-Cached	Gateway-Stored
ReadBytes	<p>The total number of bytes read from your on-premises applications in the reporting period.</p> <p>Use this metric with the <code>Sum</code> statistic to measure throughput and with the <code>Samples</code> statistic to measure IOPS.</p> <p>Units: Bytes</p>	yes	yes
ReadTime	<p>The total number of milliseconds spent to do read operations from your on-premises applications in the reporting period.</p> <p>Use this metric with the <code>Average</code> statistic to measure latency.</p> <p>Units: Milliseconds</p>	yes	yes
WriteBytes	<p>The total number of bytes written to your on-premises applications in the reporting period.</p> <p>Use this metric with the <code>Sum</code> statistic to measure throughput and with the <code>Samples</code> statistic to measure IOPS.</p> <p>Units: Bytes</p>	yes	yes
WriteTime	<p>The total number of milliseconds spent to do write operations from your on-premises applications in the reporting period.</p> <p>Use this metric with the <code>Average</code> statistic to measure latency.</p> <p>Units: Milliseconds</p>	yes	yes
QueuedWrite	<p>The number of bytes waiting to be written to AWS, sampled at the end of the reporting period.</p> <p>Units: Bytes</p>	yes	yes

Dimensions for AWS Storage Gateway Metrics

The Amazon CloudWatch namespace for the AWS Storage Gateway service is `AWS/StorageGateway`. Data is available automatically in 5-minute periods at no charge.

Dimension	Description
<code>GatewayId</code> , <code>GatewayName</code>	<p>These dimensions filter the data you request to gateway-specific metrics. You can identify a gateway to work by its <code>GatewayId</code> or its <code>GatewayName</code>. However, note that if the name of your gateway was changed for the time range that you are interested in viewing metrics, then you should use the <code>GatewayId</code>.</p> <p>Throughput and latency data of a gateway is based on all the volumes for the gateway. For information about working with gateway metrics, see Measuring Performance Between Your Gateway and AWS.</p>
<code>VolumeId</code>	<p>This dimension filters the data you request to volume-specific metrics. Identify a storage volume to work with by its <code>VolumeId</code>. For information about working with volume metrics, see Measuring Performance Between Your Application and Gateway.</p>

AWS WAF Metrics and Dimensions

AWS WAF sends data to CloudWatch every minute. For more information, see [Testing Web ACLs](#) in the *AWS WAF Developer Guide*.

AWS WAF Metrics

The `AWS/WAF` namespace includes the following metrics.

Metric	Description
<code>AllowedRequests</code>	<p>The number of allowed web requests.</p> <p>Units: Count</p> <p>Dimensions: <code>Rule</code>, <code>WebACL</code></p> <p>Valid statistics: Sum</p>
<code>BlockedRequests</code>	<p>The number of blocked web requests.</p> <p>Units: Count</p> <p>Dimensions: <code>Rule</code>, <code>WebACL</code></p> <p>Valid statistics: Sum</p>
<code>CountedRequests</code>	<p>The number of counted web requests.</p> <p>A counted web request is one that matches all of the conditions in a particular rule. Counted web requests are typically used for testing.</p> <p>Units: Count</p>

Metric	Description
	Dimensions: Rule, WebACL Valid statistics: Sum

Dimensions for AWS WAF

Dimension	Description
Rule	The name of the rule, or one of the following: <ul style="list-style-type: none"> ALL, which represents the set of all rules. Default_Action, which represents the action assigned to any request that does not match any rule with either an allow or block action.
WebACL	The name of the web ACL.

Amazon WorkSpaces Metrics and Dimensions

Amazon WorkSpaces sends data points to CloudWatch for several metrics every five minutes (five-minute metrics). Detailed monitoring, or one-minute metrics, is currently unavailable for Amazon WorkSpaces. For more information, see [Monitoring Amazon WorkSpaces](#) in the *Amazon WorkSpaces Administration Guide*.

Amazon WorkSpaces Metrics

The AWS/WorkSpaces namespace includes the following metrics.

Metric	Description	Dimensions	Statistics Available	Units
Available ¹	The number of WorkSpaces that returned a healthy status.	DirectoryId WorkspaceId	Average, Sum, Maximum, Minimum, Data Samples	Count
Unhealthy ¹	The number of WorkSpaces that returned an unhealthy status.	DirectoryId WorkspaceId	Average, Sum, Maximum, Minimum, Data Samples	Count
ConnectionAttempts ²	The number of connection attempts.	DirectoryId WorkspaceId	Average, Sum, Maximum, Minimum, Data Samples	Count
ConnectionSuccesses ²	The number of successful connections.	DirectoryId WorkspaceId	Average, Sum, Maximum, Minimum, Data Samples	Count
ConnectionFailures ²	The number of failed connections.	DirectoryId WorkspaceId	Average, Sum, Maximum,	Count

Metric	Description	Dimensions	Statistics Available	Units
			Minimum, Data Samples	
SessionLaunchTime ²	The amount of time it takes to initiate a WorkSpaces session.	DirectoryID WorkspaceID	Average, Sum, Maximum, Minimum, Data Samples	Second (time)
InSessionLatency ²	The round trip time between the WorkSpaces client and the Workspace.	DirectoryID WorkspaceID	Average, Sum, Maximum, Minimum, Data Samples	Millisecond (time)
SessionDisconnected ²	The number of connections that were closed, including user-initiated and failed connections.	DirectoryID WorkspaceID	Average, Sum, Maximum, Minimum, Data Samples	Count
UserConnected ³	The number of WorkSpaces that have a user connected.	DirectoryID WorkspaceID	Average, Sum, Maximum, Minimum, Data Samples	Count
Stopped	The number of WorkSpaces that are stopped.	DirectoryID WorkspaceID	Average, Sum, Maximum, Minimum, Data Samples	Count
Maintenance ⁴	The number of WorkSpaces that are under maintenance.	DirectoryID WorkspaceID	Average, Sum, Maximum, Minimum, Data Samples	Count

¹ Amazon WorkSpaces periodically sends status requests to a Workspace. A Workspace is marked `Available` when it responds to these requests, and `Unhealthy` when it fails to respond to these requests. These metrics are available at a per-Workspace granularity, and also aggregated for all WorkSpaces in an organization.

² Amazon WorkSpaces records metrics on connections made to each Workspace. These metrics are emitted after a user has successfully authenticated via the WorkSpaces client and the client then initiates a session. The metrics are available at a per-Workspace granularity, and also aggregated for all WorkSpaces in a directory.

³ Amazon WorkSpaces periodically sends connection status requests to a Workspace. Users are reported as connected when they are actively using their sessions. This metric is available at a per-Workspace granularity, and is also aggregated for all WorkSpaces in an organization.

⁴ This metric applies to WorkSpaces that are configured with an `AutoStop` running mode. If you have maintenance enabled for your WorkSpaces, this metric captures the number of WorkSpaces that are currently under maintenance. This metric is available at a per-Workspace granularity, which describes when a Workspace went into maintenance and when it was removed.

Dimensions for Amazon WorkSpaces Metrics

Amazon WorkSpaces metrics are available for the following dimensions.

Dimension	Description
DirectoryId	Limits the data you receive to the WorkSpaces in the specified directory. The <code>DirectoryId</code> value is in the form of <code>d-XXXXXXXXXX</code> .
WorkspaceId	Limits the data you receive to the specified WorkSpace. The <code>WorkspaceId</code> value is in the form <code>ws-XXXXXXXXXX</code> .

Creating Amazon CloudWatch Alarms


You can create a CloudWatch alarm that sends an Amazon Simple Notification Service message when the alarm changes state. An alarm watches a single metric over a time period you specify, and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The action is a notification sent to an Amazon Simple Notification Service topic or Auto Scaling policy. Alarms invoke actions for sustained state changes only. CloudWatch alarms will not invoke actions simply because they are in a particular state, the state must have changed and been maintained for a specified number of periods.

After an alarm invokes an action due to a change in state, its subsequent behavior depends on the type of action that you have associated with the alarm. For Auto Scaling policy notifications, the alarm continues to invoke the action for every period that the alarm remains in the new state. For Amazon Simple Notification Service notifications, no additional actions are invoked.

An alarm has three possible states:

- *OK*—The metric is within the defined threshold
- *ALARM*—The metric is outside of the defined threshold
- *INSUFFICIENT_DATA*—The alarm has just started, the metric is not available, or not enough data is available for the metric to determine the alarm state

In the following figure, the alarm threshold is set to 3 and the evaluation period is 3. That is, the alarm invokes its action if the oldest period is breaching and the others are breaching or missing within a time window of 3 periods. In the figure, this happens with the third through fifth time periods, and the alarm's state is set to *ALARM*. At period six, the value dips below the threshold, and the state reverts to *OK*. Later, during the ninth time period, the threshold is breached again, but the previous periods are *OK*. Consequently, the alarm's state remains *OK*.



Note

CloudWatch doesn't test or validate the actions you specify, nor does it detect any Auto Scaling or SNS errors resulting from an attempt to invoke nonexistent actions. Make sure your actions exist.

Common Features of Alarms

- You can create up to 5000 alarms per AWS account. To create or update an alarm, you use the `PutMetricAlarm` API function (`mon-put-metric-alarm` command).
- You can list any or all of the currently configured alarms, and list any alarms in a particular state using the `DescribeAlarms` API (`mon-describe-alarms` command). You can further filter the list by time range.
- You can disable and enable alarms by using the `DisableAlarmActions` and `EnableAlarmActions` APIs (`mon-disable-alarm-actions` and `mon-enable-alarm-actions` commands).
- You can test an alarm by setting it to any state using the `SetAlarmState` API (`mon-set-alarm-state` command). This temporary state change lasts only until the next alarm comparison occurs.
- You can create an alarm using the `PutMetricAlarm` API function (`mon-put-metric-alarm` command) before you've created a custom metric. In order for the alarm to be valid, you must include all of the dimensions for the custom metric in addition to the metric namespace and metric name in the alarm definition.
- Finally, you can view an alarm's history using the `DescribeAlarmHistory` API (`mon-describe-alarm-history` command). CloudWatch preserves alarm history for two weeks. Each state transition is marked with a unique time stamp. In rare cases, your history might show more than one notification for a state change. The time stamp enables you to confirm unique state changes.

Note

Some AWS resources do not send metric data to CloudWatch under certain conditions. For example, Amazon EBS may not send metric data for an available volume that is not attached to an Amazon EC2 instance, because there is no metric activity to be monitored for that volume. If you have an alarm set for such a metric, you may notice its state change to `Insufficient Data`. This may simply be an indication that your resource is inactive, and may not necessarily mean that there is a problem.

Contents

- [Set Up Amazon SNS Notifications \(p. 151\)](#)

- [Create or Edit a CloudWatch Alarm \(p. 153\)](#)
- [Create a CPU Usage Alarm that Sends Email \(p. 154\)](#)
- [Create a Load Balancer Latency Alarm that Sends Email \(p. 156\)](#)
- [Create a Storage Throughput Alarm that Sends Email \(p. 158\)](#)
- [Create Alarms to Stop, Terminate, Reboot, or Recover an Instance \(p. 159\)](#)
- [Create a Billing Alarm to Monitor Your Estimated AWS Charges \(p. 165\)](#)

Set Up Amazon SNS Notifications

Amazon CloudWatch uses Amazon Simple Notification Service (Amazon SNS) to send email. First you create and subscribe to an SNS topic. When you create a CloudWatch alarm, you can add this SNS topic to send an email notification when the alarm changes state. For more information, see the [Amazon Simple Notification Service Getting Started Guide](#).

Tip

Alternatively, if you plan to create your CloudWatch alarm using the AWS Management Console, you can skip this procedure because you can create the topic through the **Create Alarm Wizard**.

Set Up an SNS Topic Using the AWS Management Console

First you create a topic, then you subscribe to it. You can optionally publish a test message to the topic.

To create an SNS topic

1. Open the Amazon SNS console at <https://console.aws.amazon.com/sns/>.
2. On the SNS dashboard, under **Common actions**, choose **Create Topic**.
3. In the **Create new topic** dialog box, for **Topic name**, type a name for the topic (for example, my-topic).
4. Choose **Create topic**.
5. Copy the **Topic ARN** for the next task (for example, arn:aws:sns:us-east-1:111122223333:my-topic).

To subscribe to an SNS topic

1. Open the Amazon SNS console at <https://console.aws.amazon.com/sns/>.
2. In the navigation pane, choose **Subscriptions**.
3. On the **Subscriptions** page, choose **Create subscription**.
4. In the **Create subscription** dialog box, for **Topic ARN**, paste the topic ARN that you copied created in the previous task.
5. For **Protocol**, choose **Email**.
6. For **Endpoint**, type an email address that you can use to receive the notification, and then choose **Create subscription**.
7. From your email application and open the message from AWS Notifications and confirm your subscription.

Your web browser displays a confirmation response from Amazon Simple Notification Service.

To publish a test message to an SNS topic

1. Open the Amazon SNS console at <https://console.aws.amazon.com/sns/>.
2. In the navigation pane, choose **Topics**.
3. On the **Topics** page, choose the topic, and then choose **Publish to topic**.
4. In the **Publish a message** page, for **Subject**, type a subject line for your message, and for **Message**, type a brief message.
5. Choose **Publish Message**.
6. Check your email to confirm that you received the message.

Set Up an SNS Topic Using the AWS CLI

First you create an SNS topic, and then publish a message directly to the topic to test that you have properly configured it.

To set up an SNS topic

1. Create the topic using the `create-topic` command as follows.

```
aws sns create-topic --name my-topic
```

Amazon SNS returns a topic ARN with the following format:

```
{
  "TopicArn": "arn:aws:sns:us-east-1:111122223333:my-topic"
}
```

2. Subscribe your email address to the topic using the `subscribe` command. You will receive a confirmation email message if the subscription request succeeds.

```
aws sns subscribe --topic-arn arn:aws:sns:us-east-1:111122223333:my-topic
--protocol email --notification-endpoint my-email-address
```

Amazon SNS returns the following:

```
{
  "SubscriptionArn": "pending confirmation"
}
```

3. From your email application and open the message from AWS Notifications and confirm your subscription.

Your web browser displays a confirmation response from Amazon Simple Notification Service.

4. Check the subscription using the `list-subscriptions-by-topic` command.

```
aws sns list-subscriptions-by-topic --topic-arn arn:aws:sns:us-
east-1:111122223333:my-topic
```

Amazon SNS returns the following:

```
{
```

```
"Subscriptions": [
  {
    "Owner": "111122223333",
    "Endpoint": "me@mycompany.com",
    "Protocol": "email",
    "TopicArn": "arn:aws:sns:us-east-1:111122223333:my-topic",
    "SubscriptionArn": "arn:aws:sns:us-east-1:111122223333:my-
topic:64886986-bf10-48fb-a2f1-dab033aa67a3"
  }
]
```

5. (Optional) Publish a test message to the topic using the `publish` command.

```
aws sns publish --message "Verification" --topic arn:aws:sns:us-
east-1:111122223333:my-topic
```

Amazon Simple Notification Service returns the following:

```
{
  "MessageId": "42f189a0-3094-5cf6-8fd7-c2dde61a4d7d"
}
```

6. Check your email to confirm that you received the message.

Create or Edit a CloudWatch Alarm

You can choose specific metrics to trigger the alarm and specify thresholds for those metrics. You can then set your alarm to change state when a metric exceeds a threshold that you have defined. For an example of how to create an alarm that sends email, see [Creating Amazon CloudWatch Alarms](#) (p. 149).

To create an alarm

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Alarms**.
3. Choose **Create Alarm**.
4. For the **Select Metric** step, do the following:
 - a. Choose a metric category (for example, **EC2 Metrics**).
 - b. Select an instance and metric (for example, **CPUUtilization**).
 - c. For the statistic, choose one of the statistics (for example, Average) or predefined percentiles, or specify a custom percentile (for example, p95.45).
 - d. Choose a period (for example, **1 Hour**).
 - e. Choose **Next**.
5. For the **Define Alarm** step, do the following:
 - a. Under **Alarm Threshold**, type a unique name for the alarm and a description of the alarm. For **Whenever**, specify a threshold (for example, 80 percent of CPU utilization) and the number of periods.
 - b. Under **Actions**, select the type of action you want the alarm to perform when the alarm is triggered.
 - c. Choose **Create Alarm**.

To edit an alarm

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Alarms**.
3. Select the alarm, and then choose **Modify**.
4. In the **Modify Alarm** dialog box, update the alarm as necessary, and then choose **Save Changes**.

To update an email notification list that was created using the Amazon SNS console

1. Open the Amazon SNS console at <https://console.aws.amazon.com/sns/>.
2. In the navigation pane, choose **Topics**, and then choose the ARN for your notification list (topic).
3. Do one of the following:
 - To add an email address, choose **Create subscription**. For **Protocol**, choose **Email**. For **Endpoint**, type the email address of the new recipient. Choose **Create subscription**.
 - To remove an email address, choose the **Subscription ID**. Choose **Other subscription actions**, **Delete subscriptions**.
4. Choose **Publish to topic**.

Create a CPU Usage Alarm that Sends Email

You can create an Amazon CloudWatch alarm that sends an email message using Amazon Simple Notification Service (Amazon SNS) when the alarm changes state from OK to ALARM.

The alarm changes to the ALARM state when the average CPU use of an EC2 instance exceeds a specified threshold for consecutive specified periods.

Set Up a CPU Usage Alarm Using the AWS Management Console

To create an alarm that sends email based on CPU usage

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Alarms**.
3. Choose **Create Alarm**.
4. Under **EC2 Metrics**, select a metric category (for example, **Per-Instance Metrics**).
5. Select a metric as follows:
 - a. Select a row with the instance and the **CPUUtilization** metric.
 - b. For the statistic, choose **Average**, choose one of the predefined percentiles, or specify a custom percentile (for example, p95.45).
 - c. Choose a period (for example, 5 **minutes**).
 - d. Choose **Next**.

Create Alarm [X]

1. **Select Metric** 2. Define Alarm

EC2 [Search Metrics] [X] [1 to 50 of 68 Metrics]

Per-Instance Metrics [x] By Auto Scaling Group By Image (AMI) Id Aggregated by Instance Type Across All Instances

EC2 > Per-Instance Metrics

InstanceId	InstanceName	Metric Name
<input type="checkbox"/> i-0332c3c79f97a3e63		CPUCreditBalance
<input type="checkbox"/> i-0332c3c79f97a3e63		CPUCreditUsage
<input checked="" type="checkbox"/> i-0332c3c79f97a3e63		CPUUtilization
<input type="checkbox"/> i-0332c3c79f97a3e63		DiskReadBytes
<input type="checkbox"/> i-0332c3c79f97a3e63		DiskReadOps

Title: CPUUtilization [Average] [5 Minutes] [Update Graph]

Time Range: Relative Absolute UTC (GMT)
From: 3 days ago
To: 0 hours ago
Zoom: 1h | 3h | 6h | 12h | 1d | 3d | 1w | 2w

Left Y-axis: Limits Min 0 Max
Auto Auto

[Cancel] [Previous] [Next] [Create Alarm]

6. Define the alarm as follows:

- Under **Alarm Threshold**, type a unique name for the alarm (for example, myHighCpuAlarm) and a description of the alarm (for example, CPU usage exceeds 70 percent).
- Under **Whenever**, for **is**, choose **>** and type **70**. For **for**, type **2**.

Alarm Threshold

Provide the details and threshold for your alarm. Use the graph on the right to help set the appropriate threshold.

Name:

Description:

Whenever: CPUUtilization

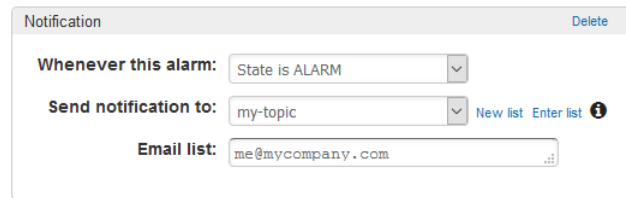
is:

for: consecutive period(s)

- Under **Actions**, for **Whenever this alarm**, select **State is ALARM**. For **Send notification to**, select an existing SNS topic or create a new one.

Actions

Define what actions are taken when your alarm changes state.



- d. To create a new SNS topic, choose **New list**. For **Send notification to**, type a name for the SNS topic (for example, myHighCpuAlarm), and for **Email list**, type a comma-separated list of email addresses to be notified when the alarm changes to the `ALARM` state. Each email address will be sent a topic subscription confirmation email. You must confirm the subscription before notifications can be sent.
- e. Choose **Create Alarm**.

Set Up a CPU Usage Alarm Using the AWS CLI

To create an alarm that sends email based on CPU usage

1. Set up an SNS topic. For more information, see [Set Up Amazon SNS Notifications \(p. 151\)](#).
2. Create an alarm using the `put-metric-alarm` command as follows.

```
aws cloudwatch put-metric-alarm --alarm-name cpu-mon --alarm-  
description "Alarm when CPU exceeds 70%" --metric-name CPUUtilization  
--namespace AWS/EC2 --statistic Average --period 300 --threshold 70  
--comparison-operator GreaterThanThreshold --dimensions  
Name=InstanceId,Value=i-12345678 --evaluation-periods 2 --alarm-actions  
arn:aws:sns:us-east-1:111122223333:my-topic --unit Percent
```

3. Test the alarm by forcing an alarm state change using the `set-alarm-state` command.
 - a. Change the alarm state from `INSUFFICIENT_DATA` to `OK`:

```
aws cloudwatch set-alarm-state --alarm-name cpu-mon --state-reason  
"initializing" --state-value OK
```

- b. Change the alarm state from `OK` to `ALARM`:

```
aws cloudwatch set-alarm-state --alarm-name cpu-mon --state-reason  
"initializing" --state-value ALARM
```

- c. Check that you have received an email notification about the alarm.

Create a Load Balancer Latency Alarm that Sends Email

You can set up an Amazon SNS notification and configure an alarm that monitors latency exceeding 100 ms for your Classic Load Balancer.

Set Up a Latency Alarm Using the AWS Management Console

To create a load balancer latency alarm that sends email

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Alarms**.
3. Choose **Create Alarm**.
4. Under **CloudWatch Metrics by Category**, select the **ELB Metrics** category.
5. Select the row with the Classic Load Balancer and the **Latency** metric.
6. For the statistic, choose **Average**, choose one of the predefined percentiles, or specify a custom percentile (for example, p95.45).
7. For the period, choose **1 Minute**.
8. Choose **Next**.
9. Under **Alarm Threshold**, type a unique name for the alarm (for example, `myHighCpuAlarm`) and a description of the alarm (for example, Alarm when Latency exceeds 100s).
10. Under **Whenever**, for **is**, select **>** and type **0.1**. For **for**, type **3**.
11. Under **Actions**, for **Whenever this alarm**, choose **State is ALARM**. For **Send notification to** choose an existing SNS topic or create a new one.

To create an SNS topic, choose **New list**. For **Send notification to**, type a name for the SNS topic (for example, `myHighCpuAlarm`), and for **Email list**, type a comma-separated list of email addresses to be notified when the alarm changes to the `ALARM` state. Each email address will be sent a topic subscription confirmation email. You must confirm the subscription before notifications can be sent.

12. Choose **Create Alarm**.

Set Up a Latency Alarm Using the AWS CLI

To create a load balancer latency alarm that sends email

1. Set up an SNS topic. For more information, see [Set Up Amazon SNS Notifications \(p. 151\)](#)
2. Create the alarm using the `put-metric-alarm` command as follows:

```
aws cloudwatch put-metric-alarm --alarm-name lb-mon --alarm-description  
"Alarm when Latency exceeds 100s" --metric-name Latency --namespace  
AWS/ELB --statistic Average --period 60 --threshold 100 --comparison-  
operator GreaterThanThreshold --dimensions Name=LoadBalancerName,Value=my-  
server --evaluation-periods 3 --alarm-actions arn:aws:sns:us-  
east-1:111122223333:my-topic --unit Seconds
```

3. Test the alarm by forcing an alarm state change using the `set-alarm-state` command.
 - a. Change the alarm state from `INSUFFICIENT_DATA` to `OK`:

```
aws cloudwatch set-alarm-state --alarm-name lb-mon --state-reason  
"initializing" --state-value OK
```

- b. Change the alarm state from `OK` to `ALARM`:

```
aws cloudwatch set-alarm-state --alarm-name lb-mon --state-reason  
"initializing" --state-value ALARM
```

- c. Check that you have received an email notification about the alarm.

Create a Storage Throughput Alarm that Sends Email

You can set up an SNS notification and configure an alarm that sends email when Amazon EBS exceeds 100 MB throughput.

Set Up a Storage Throughput Alarm Using the AWS Management Console

To create a storage throughput alarm that sends email

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Alarms**.
3. Choose **Create Alarm**.
4. Under **EBS Metrics**, choose a metric category.
5. Select the row with the volume and the **VolumeWriteBytes** metric.
6. For the statistic, choose **Average**.
7. For the period, choose **5 Minutes**.
8. Choose **Next**.
9. Under **Alarm Threshold**, type a unique name for the alarm (for example, myHighWriteAlarm) and a description of the alarm (for example, VolumeWriteBytes exceeds 100,000 KiB/s).
10. Under **Whenever**, for **is**, choose **>** and type 100000. For **for**, type 15 consecutive periods.

A graphical representation of the threshold is shown under **Alarm Preview**.

11. Under **Actions**, for **Whenever this alarm**, choose **State is ALARM**. For **Send notification to**, chose an existing SNS topic or create one.

To create an SNS topic, choose **New list**. For **Send notification to**, type a name for the SNS topic (for example, myHighCpuAlarm), and for **Email list**, type a comma-separated list of email addresses to be notified when the alarm changes to the `ALARM` state. Each email address will be sent a topic subscription confirmation email. You must confirm the subscription before notifications can be sent to an email address.

12. Choose **Create Alarm**.

Set Up a Storage Throughput Alarm Using the AWS CLI

To create a storage throughput alarm that sends email

1. Create an SNS topic. For more information, see [Set Up Amazon SNS Notifications](#) (p. 151).
2. Create the alarm.

```
aws cloudwatch put-metric-alarm --alarm-name ebs-mon --alarm-description  
"Alarm when EBS volume exceeds 100MB throughput" --metric-name  
VolumeReadBytes --namespace AWS/EBS --statistic Average --period 300  
--threshold 100000000 --comparison-operator GreaterThanThreshold --  
dimensions Name=VolumeId,Value=my-volume-id --evaluation-periods 3  
--alarm-actions arn:aws:sns:us-east-1:111122223333:my-alarm-topic  
--insufficient-data-actions arn:aws:sns:us-east-1:111122223333:my-  
insufficient-data-topic
```

3. Test the alarm by forcing an alarm state change using the `set-alarm-state` command.
 - a. Change the alarm state from `INSUFFICIENT_DATA` to `OK`:

```
aws cloudwatch set-alarm-state --alarm-name ebs-mon --state-reason  
"initializing" --state-value OK
```

- b. Change the alarm state from `OK` to `ALARM`:

```
aws cloudwatch set-alarm-state --alarm-name ebs-mon --state-reason  
"initializing" --state-value ALARM
```

- c. Change the alarm state from `ALARM` to `INSUFFICIENT_DATA`:

```
aws cloudwatch set-alarm-state --alarm-name ebs-mon --state-reason  
"initializing" --state-value INSUFFICIENT_DATA
```

- d. Check that you have received an email notification about the alarm.

Create Alarms to Stop, Terminate, Reboot, or Recover an Instance

Using Amazon CloudWatch alarm actions, you can create alarms that automatically stop, terminate, reboot, or recover your EC2 instances. You can use the stop or terminate actions to help you save money when you no longer need an instance to be running. You can use the reboot and recover actions to automatically reboot those instances or recover them onto new hardware if a system impairment occurs.

Every alarm action you create uses alarm action ARNs. One set of ARNs is more secure because it requires you to have the `EC2ActionsAccess` IAM role in your account. This IAM role enables you to perform stop, terminate, or reboot actions—previously you could not execute an action if you were using an IAM role. Existing alarms that use the previous alarm action ARNs do not require this IAM role, however it is recommended that you change the ARN and add the role when you edit an existing alarm that uses these ARNs.

The `EC2ActionsAccess` IAM role enables AWS to perform alarm actions on your behalf. When you create an alarm action for the first time using the Amazon EC2 or Amazon CloudWatch consoles, AWS automatically creates this role for you.

There are a number of scenarios in which you might want to automatically stop or terminate your instance. For example, you might have instances dedicated to batch payroll processing jobs or scientific computing tasks that run for a period of time and then complete their work. Rather than letting those instances sit idle (and accrue charges), you can stop or terminate them which can help you to save money. The main difference between using the stop and the terminate alarm actions is that you can easily restart a stopped instance if you need to run it again later, and you can keep the same

instance ID and root volume. However, you cannot restart a terminated instance. Instead, you must launch a new instance.

You can add the stop, terminate, reboot, or recover actions to any alarm that is set on an Amazon EC2 per-instance metric, including basic and detailed monitoring metrics provided by Amazon CloudWatch (in the AWS/EC2 namespace), as well as any custom metrics that include the "InstanceId=" dimension, as long as the InstanceId value refers to a valid running Amazon EC2 instance.

Console Support

You can create alarms using the CloudWatch console or the Amazon EC2 console. The procedures in this documentation use the CloudWatch console. For procedures that use the Amazon EC2 console, see [Create Alarms That Stop, Terminate, Reboot, or Recover an Instance](#) in the *Amazon EC2 User Guide for Linux Instances*.

Permissions

If you are using an AWS Identity and Access Management (IAM) account to create or modify an alarm, you must have the following permissions:

- `ec2:DescribeInstanceStatus` and `ec2:DescribeInstances` — For all alarms on Amazon EC2 instance status metrics
- `ec2:StopInstances` — For alarms with stop actions
- `ec2:TerminateInstances` — For alarms with terminate actions
- `ec2:DescribeInstanceRecoveryAttribute` and `ec2:RecoverInstances` — For alarms with recover actions

If you have read/write permissions for Amazon CloudWatch but not for Amazon EC2, you can still create an alarm but the stop or terminate actions won't be performed on the instance. However, if you are later granted permission to use the associated Amazon EC2 APIs, the alarm actions you created earlier will be performed. For more information, see [Permissions and Policies](#) in the *IAM User Guide*.

If you want to use an IAM role to stop, terminate, or reboot an instance using an alarm action, you can only use the `EC2ActionsAccess` role. Other IAM roles are not supported. If you are using another IAM role, you cannot stop, terminate, or reboot the instance. However, you can still see the alarm state and perform any other actions such as Amazon SNS notifications or Auto Scaling policies.

If you are using temporary security credentials granted using the AWS Security Token Service (AWS STS), you cannot recover an Amazon EC2 instance using alarm actions.

Contents

- [Adding Stop Actions to Amazon CloudWatch Alarms](#) (p. 160)
- [Adding Terminate Actions to Amazon CloudWatch Alarms](#) (p. 161)
- [Adding Reboot Actions to Amazon CloudWatch Alarms](#) (p. 162)
- [Adding Recover Actions to Amazon CloudWatch Alarms](#) (p. 163)
- [Viewing the History of Triggered Alarms and Actions](#) (p. 165)

Adding Stop Actions to Amazon CloudWatch Alarms

You can create an alarm that stops an Amazon EC2 instance when a certain threshold has been met. For example, you may run development or test instances and occasionally forget to shut them off. You can create an alarm that is triggered when the average CPU utilization percentage has been lower than 10 percent for 24 hours, signaling that it is idle and no longer in use. You can adjust the threshold, duration, and period to suit your needs, plus you can add an SNS notification, so that you will receive an email when the alarm is triggered.

Amazon EC2 instances that use an Amazon Elastic Block Store volume as the root device can be stopped or terminated, whereas instances that use the instance store as the root device can only be terminated.

To create an alarm to stop an idle instance using the Amazon CloudWatch console

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Alarms**.
3. Choose **Create Alarm**.
4. For the **Select Metric** step, do the following:
 - a. Under **EC2 Metrics**, choose **Per-Instance Metrics**.
 - b. Select the row with the instance and the **CPUUtilization** metric.
 - c. For the statistic, choose **Average**.
 - d. Choose a period (for example, 1 **Hour**).
 - e. Choose **Next**.
5. For the **Define Alarm** step, do the following:
 - a. Under **Alarm Threshold**, type a unique name for the alarm (for example, Stop EC2 instance) and a description of the alarm (for example, Stop EC2 instance when CPU is idle too long).
 - b. Under **Whenever**, for **is**, choose **<** and type 10. For **for**, type 24 consecutive periods.

A graphical representation of the threshold is shown under **Alarm Preview**.
 - c. Under **Notification**, for **Send notification to**, choose an existing SNS topic or create a new one.

To create an SNS topic, choose **New list**. For **Send notification to**, type a name for the SNS topic (for example, Stop_EC2_Instance), and for **Email list**, type a comma-separated list of email addresses to be notified when the alarm changes to the `ALARM` state. Each email address will be sent a topic subscription confirmation email. You must confirm the subscription before notifications can be sent to an email address.
 - d. Choose **EC2 Action**.
 - e. Under **EC2 Action**, for **Whenever this alarm**, choose **State is ALARM**. For **Take this action**, choose **Stop this instance**.
 - f. If prompted, select **Create IAM role: EC2ActionsAccess** to automatically create an IAM role so that AWS can automatically stop the instance on your behalf when the alarm is triggered.
 - g. Choose **Create Alarm**.

Adding Terminate Actions to Amazon CloudWatch Alarms

You can create an alarm that terminates an EC2 instance automatically when a certain threshold has been met (as long as termination protection is not enabled for the instance). For example, you might want to terminate an instance when it has completed its work, and you don't need the instance again. If you might want to use the instance later, you should stop the instance instead of terminating it. For information about enabling and disabling termination protection for an instance, see [Enabling Termination Protection for an Instance](#) in the *Amazon EC2 User Guide for Linux Instances*.

To create an alarm to terminate an idle instance using the Amazon CloudWatch console

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Alarms**.
3. Choose **Create Alarm**.

4. For the **Select Metric** step, do the following:
 - a. Under **EC2 Metrics**, choose **Per-Instance Metrics**.
 - b. Select the row with the instance and the **CPUUtilization** metric.
 - c. For the statistic, choose **Average**.
 - d. Choose a period (for example, 1 **Hour**).
 - e. Choose **Next**.
5. For the **Define Alarm** step, do the following:
 - a. Under **Alarm Threshold**, type a unique name for the alarm (for example, Terminate EC2 instance) and a description of the alarm (for example, Terminate EC2 instance when CPU is idle for too long).
 - b. Under **Whenever**, for **is**, choose **<** and type 10. For **for**, type 24 consecutive periods.

A graphical representation of the threshold is shown under **Alarm Preview**.

- c. Under **Notification**, for **Send notification to**, choose an existing SNS topic or create a new one.

To create an SNS topic, choose **New list**. For **Send notification to**, type a name for the SNS topic (for example, Terminate_EC2_Instance), and for **Email list**, type a comma-separated list of email addresses to be notified when the alarm changes to the `ALARM` state. Each email address will be sent a topic subscription confirmation email. You must confirm the subscription before notifications can be sent to an email address.

- d. Choose **EC2 Action**.
- e. Under **EC2 Action**, for **Whenever this alarm**, choose **State is ALARM**. For **Take this action**, choose **Terminate this instance**.
- f. If prompted, select **Create IAM role: EC2ActionsAccess** to automatically create an IAM role so that AWS can automatically terminate the instance on your behalf when the alarm is triggered.
- g. Choose **Create Alarm**.

Adding Reboot Actions to Amazon CloudWatch Alarms

You can create an Amazon CloudWatch alarm that monitors an Amazon EC2 instance and automatically reboots the instance. The reboot alarm action is recommended for Instance Health Check failures (as opposed to the recover alarm action, which is suited for System Health Check failures). An instance reboot is equivalent to an operating system reboot. In most cases, it takes only a few minutes to reboot your instance. When you reboot an instance, it remains on the same physical host, so your instance keeps its public DNS name, private IP address, and any data on its instance store volumes.

Rebooting an instance doesn't start a new instance billing hour, unlike stopping and restarting your instance. For more information about rebooting an instance, see [Reboot Your Instance](#) in the *Amazon EC2 User Guide for Linux Instances*.

Important

To avoid a race condition between the reboot and recover actions, we recommend that you set the alarm threshold to **3** for **1** minute when creating alarms that reboot an Amazon EC2 instance.

To create an alarm to reboot an instance using the Amazon CloudWatch console

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.

2. In the navigation pane, choose **Alarms**.
3. Choose **Create Alarm**.
4. For the **Select Metric** step, do the following:
 - a. Under **EC2 Metrics**, choose **Per-Instance Metrics**.
 - b. Select the row with the instance and the **StatusCheckFailed_Instance** metric.
 - c. For the statistic, choose **Minimum**.
 - d. Choose a period (for example, **1 minute**).
 - e. Choose **Next**.
5. For the **Define Alarm** step, do the following:
 - a. Under **Alarm Threshold**, type a unique name for the alarm (for example, Reboot EC2 instance) and a description of the alarm (for example, Reboot EC2 instance when health checks fail).
 - b. Under **Whenever**, for **is**, choose **>** and type 0. For **for**, type 3 consecutive periods.

A graphical representation of the threshold is shown under **Alarm Preview**.
 - c. Under **Notification**, for **Send notification to**, choose an existing SNS topic or create a new one.

To create an SNS topic, choose **New list**. For **Send notification to**, type a name for the SNS topic (for example, Reboot_EC2_Instance), and for **Email list**, type a comma-separated list of email addresses to be notified when the alarm changes to the `ALARM` state. Each email address will be sent a topic subscription confirmation email. You must confirm the subscription before notifications can be sent to an email address.
 - d. Choose **EC2 Action**.
 - e. Under **EC2 Action**, for **Whenever this alarm**, choose **State is ALARM**. For **Take this action**, choose **Reboot this instance**.
 - f. If prompted, select **Create IAM role: EC2ActionsAccess** to automatically create an IAM role so that AWS can automatically stop the instance on your behalf when the alarm is triggered.
 - g. Choose **Create Alarm**.

Adding Recover Actions to Amazon CloudWatch Alarms

You can create an Amazon CloudWatch alarm that monitors an Amazon EC2 instance and automatically recovers the instance if it becomes impaired due to an underlying hardware failure or a problem that requires AWS involvement to repair. Terminated instances cannot be recovered. A recovered instance is identical to the original instance, including the instance ID, private IP addresses, Elastic IP addresses, and all instance metadata.

When the `StatusCheckFailed_System` alarm is triggered, and the recover action is initiated, you will be notified by the Amazon SNS topic that you chose when you created the alarm and associated the recover action. During instance recovery, the instance is migrated during an instance reboot, and any data that is in-memory is lost. When the process is complete, information is published to the SNS topic you've configured for the alarm. Anyone who is subscribed to this SNS topic will receive an email notification that includes the status of the recovery attempt and any further instructions. You will notice an instance reboot on the recovered instance.

Examples of problems that cause system status checks to fail include:

- Loss of network connectivity
- Loss of system power

- Software issues on the physical host
- Hardware issues on the physical host

The recover action is only supported on:

- The C3, C4, M3, M4, R3, T2, and X1 instance types
- Instances in a VPC
- Instances with shared tenancy (the tenancy attribute is set to `default`)
- Instances that use Amazon EBS storage exclusively

If your instance has a public IP address, it will retain the same public IP address after recovery.

To create an alarm to recover an instance using the Amazon CloudWatch console

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Alarms**.
3. Choose **Create Alarm**.
4. For the **Select Metric** step, do the following:
 - a. Under **EC2 Metrics**, choose **Per-Instance Metrics**.
 - b. Select the row with the instance and the **StatusCheckFailed_System** metric.
 - c. For the statistic, choose **Minimum**.
 - d. Choose a period (for example, **1 Minute**).
5. For the **Define Alarm** step, do the following:
 - a. Under **Alarm Threshold**, type a unique name for the alarm (for example, Recover EC2 instance) and a description of the alarm (for example, Recover EC2 instance when health checks fail).
 - b. Under **Whenever**, for **is**, choose **>** and type 0. For **for**, type 2 consecutive periods.

Important

To avoid a race condition between the reboot and recover actions, we recommend that you set the alarm threshold to **2** for **1 Minute** when creating alarms that recover an EC2 instance.

- c. Under **Notification**, for **Send notification to**, choose an existing SNS topic or create a new one.

To create an SNS topic, choose **New list**. For **Send notification to**, type a name for the SNS topic (for example, Recover_EC2_Instance), and for **Email list**, type a comma-separated list of email addresses to be notified when the alarm changes to the `ALARM` state. Each email address will be sent a topic subscription confirmation email. You must confirm the subscription before notifications can be sent to an email address.
- d. Choose **EC2 Action**.
- e. Under **EC2 Action**, for **Whenever this alarm**, choose **State is ALARM**. For **Take this action**, choose **Recover this instance**.
- f. If prompted, select **Create IAM role: EC2ActionsAccess** to automatically create an IAM role so that AWS can automatically stop the instance on your behalf when the alarm is triggered.
- g. Choose **Create Alarm**.

Viewing the History of Triggered Alarms and Actions

You can view alarm and action history in the Amazon CloudWatch console. Amazon CloudWatch keeps the last two weeks' worth of alarm and action history.

To view the history of triggered alarms and actions

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Alarms**.
3. Choose the alarm.
4. Choose the **Details** tab to view the most recent state transition along with the time and metric values.
5. Choose the **History** tab to view the most recent history entries.

Create a Billing Alarm to Monitor Your Estimated AWS Charges

You can monitor your estimated AWS charges using Amazon CloudWatch. When you enable the monitoring of estimated charges for your AWS account, the estimated charges are calculated and sent several times daily to CloudWatch as metric data.

Billing metric data is stored in the US East (N. Virginia) region and represents worldwide charges. This data includes the estimated charges for every service in AWS that you use, as well as the estimated overall total of your AWS charges.

You can choose to receive alerts by email when charges have exceeded a certain threshold. These alerts are triggered by CloudWatch and messages are sent using Amazon Simple Notification Service (Amazon SNS).

Tasks

- [Enable Billing Alerts \(p. 165\)](#)
- [Create a Billing Alarm \(p. 166\)](#)
- [Check the Alarm Status \(p. 167\)](#)
- [Delete a Billing Alarm \(p. 168\)](#)

Enable Billing Alerts

Before you can create an alarm for your estimated charges, you must enable billing alerts, so that you can monitor your estimated AWS charges and create an alarm using billing metric data. After you enable billing alerts, you cannot disable data collection, but you can delete any billing alarms you created.

After you enable billing alerts for the first time, it takes about 15 minutes before you can view billing data and set billing alarms.

Requirements

- You must be signed in using root account credentials; IAM users cannot enable billing alerts for your AWS account.

- For consolidated billing accounts, billing data for each linked account can be found by logging in as the paying account. You can view billing data for total estimated charges and estimated charges by service for each linked account as well as for the consolidated account.

To enable the monitoring of estimated charges

1. Open the Billing and Cost Management console at <https://console.aws.amazon.com/billing/home?#>.
2. In the navigation pane, choose **Preferences**.
3. Select **Receive Billing Alerts**.

The screenshot shows the 'Preferences' page in the AWS Billing and Cost Management console. On the left is a navigation menu with items like Dashboard, Bills, Cost Explorer, Budgets, Reports, Cost Allocation Tags, Payment Methods, Payment History, Consolidated Billing, Preferences (highlighted), Credits, Tax Settings, and DevPay. The main content area has the title 'Preferences' and a help icon. It contains three sections: 1. 'Receive PDF Invoice By Email' with an unchecked checkbox and a description: 'Turn on this feature to receive a PDF version of your invoice by email. Invoices are generally available within the first three days of the month.' 2. 'Receive Billing Alerts' with a checked checkbox and a description: 'Turn on this feature to monitor your AWS usage charges and recurring fees automatically, making it easier to track and manage your spending on AWS. You can set up billing alerts to receive email notifications when your charges reach a specified threshold. Once enabled, this preference cannot be disabled. [Manage Billing Alerts](#)' 3. 'Receive Billing Reports' with an unchecked checkbox and a description: 'Turn on this feature to receive ongoing reports of your AWS charges once or more daily. AWS delivers these reports to the Amazon S3 bucket that you specify where indicated below. For consolidated billing customers, AWS generates reports only for paying accounts. Linked accounts cannot sign up for billing reports.' Below this is a 'Save to S3 Bucket:' label, a text input field containing 'bucket name', and a 'Verify' button. At the bottom is a blue 'Save preferences' button.

4. Choose **Save preferences**.

Create a Billing Alarm

After you've enabled billing alerts, you can create a billing alarm. In this procedure, you create an alarm that sends an email message when your estimated charges for AWS exceed a specified threshold.

Note

This procedure uses the advanced options. To use the simple options, see [Create a Billing Alarm](#) (p. 13) in *Monitor Your Estimated Charges Using CloudWatch*.

To create a billing alarm using the Amazon CloudWatch console

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. If necessary, change the region to US East (N. Virginia). Billing metric data is stored in this region and represents worldwide charges.
3. In the navigation pane, choose **Alarms, Billing**.
4. Choose **Create Alarm**.
5. Choose **show advanced** to switch to the advanced options.
6. Under **Alarm Threshold**, replace the default name for the alarm (for example, My Estimated Charges) and a description for the alarm (for example, Estimated Monthly Charges).
7. Under **Whenever charges for**, for **is**, choose **>=** and then type the monetary amount (for example, 200) that must be exceeded to trigger the alarm and send an email.

Tip

Under **Alarm Preview**, there is an estimate of your charges that you can use to set an appropriate amount.

Modify Alarm

Alarm Threshold

Provide the details and threshold for your alarm. Use the graph on the right to help set the appropriate threshold.

Name: My Estimated Charges

Description: Estimated Monthly Charges

Whenever charges for: EstimatedCharges

is: >= USD \$ 200

Alarm Preview

This alarm will trigger when the blue line goes up to or above the red line

EstimatedCharges >= 200

250
200
150
100
50
0

10/25 00:00 10/27 00:00 10/29 00:00

Namespace: AWS/Billing

Currency: USD

Metric Name: EstimatedCharges

Actions

Define what actions are taken when your alarm changes state.

Notification Delete

Whenever this alarm: State is ALARM

Send notification to: Select a notification list New list Enter list i

+ Notification + AutoScaling Action + EC2 Action

[show simple](#) | showing advanced options
Showing simple options will revert any changes you have made above.

Cancel Previous Next **Create Alarm**

- Under **Actions**, for **Whenever this alarm**, choose **State is ALARM**. For **Send notification to**, choose an existing SNS topic or create a new one.

To create an SNS topic, choose **New list**. For **Send notification to**, type a comma-separated list of SNS topics, and for **Email list** box, type a comma-separated list of email addresses where email notifications should be sent. Each email address will be sent a topic subscription confirmation email. You must confirm the subscription before notifications can be sent to an email address.

- Choose **Create Alarm**.

Check the Alarm Status

You can check the status of your billing alarm.

To check alarm status

- Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
- If necessary, change the region to US East (N. Virginia). Billing metric data is stored in this region and reflects worldwide charges.
- In the navigation pane, choose **Alarms, Billing**.
- Select the check box next to the alarm. Note that until the subscription is confirmed, it is shown as "Pending confirmation". After the subscription is confirmed, refresh the console to show the updated status.

Delete a Billing Alarm

You can delete your billing alarm when you no longer need it.

To delete a billing alarm

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. If necessary, change the region to US East (N. Virginia). Billing metric data is stored in this region and reflects worldwide charges.
3. In the navigation pane, choose **Alarms, Billing**.
4. Select the check box next to the alarm and then choose **Delete**.
5. When prompted for confirmation, choose **Yes, Delete**.

Authentication and Access Control for Amazon CloudWatch

Access to Amazon CloudWatch requires credentials. Those credentials must have permissions to access AWS resources, such as retrieving CloudWatch metric data about your cloud resources. The following sections provide details about how you can use [AWS Identity and Access Management \(IAM\)](#) and CloudWatch to help secure your resources by controlling who can access them:

- [Authentication](#) (p. 169)
- [Access Control](#) (p. 170)

Authentication

You can access AWS as any of the following types of identities:

- **AWS account root user** – When you sign up for AWS, you provide an email address and password that is associated with your AWS account. These are your *root credentials* and they provide complete access to all of your AWS resources.

Important

For security reasons, we recommend that you use the root credentials only to create an *administrator user*, which is an *IAM user* with full permissions to your AWS account. Then, you can use this administrator user to create other IAM users and roles with limited permissions. For more information, see [IAM Best Practices](#) and [Creating an Admin User and Group](#) in the *IAM User Guide*.

- **IAM user** – An *IAM user* is simply an identity within your AWS account that has specific custom permissions (for example, permissions to view metrics in CloudWatch). You can use an IAM user name and password to sign in to secure AWS webpages like the [AWS Management Console](#), [AWS Discussion Forums](#), or the [AWS Support Center](#).

In addition to a user name and password, you can also generate [access keys](#) for each user. You can use these keys when you access AWS services programmatically, either through [one of the several SDKs](#) or by using the [AWS Command Line Interface \(CLI\)](#). The SDK and CLI tools use the access keys to cryptographically sign your request. If you don't use the AWS tools, you must sign the request yourself. CloudWatch supports *Signature Version 4*, a protocol for authenticating inbound API requests. For more information about authenticating requests, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

- **IAM role** – An [IAM role](#) is another IAM identity you can create in your account that has specific permissions. It is similar to an *IAM user*, but it is not associated with a specific person. An IAM role enables you to obtain temporary access keys that can be used to access AWS services and resources. IAM roles with temporary credentials are useful in the following situations:
 - **Federated user access** – Instead of creating an IAM user, you can use preexisting user identities from AWS Directory Service, your enterprise user directory, or a web identity provider. These are known as *federated users*. AWS assigns a role to a federated user when access is requested through an [identity provider](#). For more information about federated users, see [Federated Users and Roles](#) in the *IAM User Guide*.
 - **Cross-account access** – You can use an IAM role in your account to grant another AWS account permissions to access your account's resources. For an example, see [Tutorial: Delegate Access Across AWS Accounts Using IAM Roles](#) in the *IAM User Guide*.
 - **AWS service access** – You can use an IAM role in your account to grant an AWS service permissions to access your account's resources. For example, you can create a role that allows Amazon Redshift to access an Amazon S3 bucket on your behalf and then load data stored in the bucket into an Amazon Redshift cluster. For more information, see [Creating a Role to Delegate Permissions to an AWS Service](#) in the *IAM User Guide*.
 - **Applications running on Amazon EC2** – Instead of storing access keys within the EC2 instance for use by applications running on the instance and making API requests, you can use an IAM role to manage temporary credentials for these applications. To assign an AWS role to an EC2 instance and make it available to all of its applications, you can create an instance profile that is attached to the instance. An instance profile contains the role and enables programs running on the EC2 instance to get temporary credentials. For more information, see [Using Roles for Applications on Amazon EC2](#) in the *IAM User Guide*.

Access Control

You can have valid credentials to authenticate your requests, but unless you have permissions you cannot create or access CloudWatch resources. For example, you must have permissions to create CloudWatch dashboard widgets, view metrics, and so on.

The following sections describe how to manage permissions for CloudWatch. We recommend that you read the overview first.

- [Overview of Managing Access Permissions to Your CloudWatch Resources](#) (p. 171)
- [Using Identity-Based Policies \(IAM Policies\) for CloudWatch](#) (p. 175)
- [Amazon CloudWatch Permissions Reference](#) (p. 180)

Overview of Managing Access Permissions to Your CloudWatch Resources

Every AWS resource is owned by an AWS account, and permissions to create or access a resource are governed by permissions policies. An account administrator can attach permissions policies to IAM identities (that is, users, groups, and roles), and some services (such as AWS Lambda) also support attaching permissions policies to resources.

Note

An *account administrator* (or administrator user) is a user with administrator privileges. For more information, see [IAM Best Practices](#) in the *IAM User Guide*.

When granting permissions, you decide who is getting the permissions, the resources they get permissions for, and the specific actions that you want to allow on those resources.

Topics

- [CloudWatch Resources and Operations](#) (p. 171)
- [Understanding Resource Ownership](#) (p. 173)
- [Managing Access to Resources](#) (p. 173)
- [Specifying Policy Elements: Actions, Effects, and Principals](#) (p. 174)
- [Specifying Conditions in a Policy](#) (p. 174)

CloudWatch Resources and Operations

CloudWatch doesn't have any specific resources for you to control access to. Therefore, there are no CloudWatch Amazon Resource Names (ARNs) for you to use in an IAM policy. For example, you can't give a user access to CloudWatch data for only a specific set of EC2 instances or a specific load balancer or. Permissions granted using IAM cover all the cloud resources you use or monitor with CloudWatch. In addition, you can't use IAM roles with the CloudWatch command line tools.

You use an * (asterisk) as the resource when writing a policy to control access to CloudWatch actions. For example:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": ["cloudwatch:GetMetricStatistics", "cloudwatch:ListMetrics"],
    "Resource": "*",
    "Condition": {
      "Bool": {
        "aws:SecureTransport": "true"
      }
    }
  }]
}
```

For more information about ARNs, see [ARNs](#) in *IAM User Guide*. For information about CloudWatch Logs ARNs, see [Amazon Resource Names \(ARNs\) and AWS Service Namespaces](#) in the *Amazon Web Services General Reference*. For an example of a policy that covers CloudWatch actions, see [Using Identity-Based Policies \(IAM Policies\) for CloudWatch](#) (p. 175).

Action	ARN (with region)	ARN (for use with IAM role)
<i>Stop</i>	arn:aws:automate:us-east-1:ec2:stop	arn:aws:swf:us-east-1: <i>customer-account</i> :action/actions/AWS_EC2.InstanceId.Stop/1.0 Note You must create at least one stop alarm using the Amazon EC2 or CloudWatch console to create the EC2ActionsAccess IAM role. After this IAM role is created, you can create stop alarms using the CLI.
<i>Terminate</i>	arn:aws:automate:us-east-1:ec2:terminate	arn:aws:swf:us-east-1: <i>customer-account</i> :action/actions/AWS_EC2.InstanceId.Terminate/1.0 Note You must create at least one terminate alarm using the Amazon EC2 or CloudWatch console to create the EC2ActionsAccess IAM role. After this IAM role is created, you can create terminate alarms using the CLI.
<i>Reboot</i>	n/a	arn:aws:swf:us-east-1: <i>customer-account</i> :action/actions/AWS_EC2.InstanceId.Reboot/1.0 Note You must create at least one reboot alarm using the Amazon EC2 or CloudWatch console to create the EC2ActionsAccess IAM role. After this IAM role is created, you can create reboot alarms using the CLI.
<i>Recover</i>	arn:aws:automate:us-east-1:ec2:recover	n/a

Understanding Resource Ownership

The AWS account owns the resources that are created in the account, regardless of who created the resources. Specifically, the resource owner is the AWS account of the [principal entity](#) (that is, the root account, an IAM user, or an IAM role) that authenticates the resource creation request. CloudWatch does not have any resources that you can own.

Managing Access to Resources

A *permissions policy* describes who has access to what. The following section explains the available options for creating permissions policies.

Note

This section discusses using IAM in the context of CloudWatch. It doesn't provide detailed information about the IAM service. For complete IAM documentation, see [What Is IAM?](#) in the *IAM User Guide*. For information about IAM policy syntax and descriptions, see [AWS IAM Policy Reference](#) in the *IAM User Guide*.

Policies attached to an IAM identity are referred to as identity-based policies (IAM policies) and policies attached to a resource are referred to as resource-based policies. CloudWatch supports only identity-based policies (IAM policies).

Topics

- [Identity-Based Policies \(IAM Policies\) \(p. 173\)](#)
- [Resource-Based Policies \(p. 174\)](#)

Identity-Based Policies (IAM Policies)

You can attach policies to IAM identities. For example, you can do the following:

- **Attach a permissions policy to a user or a group in your account** – To grant a user permissions to create an Amazon CloudWatch resource, such as metrics, you can attach a permissions policy to a user or group that the user belongs to.
- **Attach a permissions policy to a role (grant cross-account permissions)** – You can attach an identity-based permissions policy to an IAM role to grant cross-account permissions. For example, the administrator in account A can create a role to grant cross-account permissions to another AWS account (for example, account B) or an AWS service as follows:
 1. Account A administrator creates an IAM role and attaches a permissions policy to the role that grants permissions on resources in account A.
 2. Account A administrator attaches a trust policy to the role identifying account B as the principal who can assume the role.
 3. Account B administrator can then delegate permissions to assume the role to any users in account B. Doing this allows users in account B to create or access resources in account A. The principal in the trust policy can also be an AWS service principal if you want to grant an AWS service permissions to assume the role.

For more information about using IAM to delegate permissions, see [Access Management](#) in the *IAM User Guide*.

For more information about using identity-based policies with CloudWatch, see [Using Identity-Based Policies \(IAM Policies\) for CloudWatch \(p. 175\)](#). For more information about users, groups, roles, and permissions, see [Identities \(Users, Groups, and Roles\)](#) in the *IAM User Guide*.

Resource-Based Policies

Other services, such as Amazon S3, also support resource-based permissions policies. For example, you can attach a policy to an Amazon S3 bucket to manage access permissions to that bucket. CloudWatch doesn't support resource-based policies.

Specifying Policy Elements: Actions, Effects, and Principals

For each CloudWatch resource, the service defines a set of API operations. To grant permissions for these API operations, CloudWatch defines a set of actions that you can specify in a policy. Some API operations can require permissions for more than one action in order to perform the API operation. For more information about resources and API operations, see [CloudWatch Resources and Operations \(p. 171\)](#) and [CloudWatch Actions](#).

The following are the basic policy elements:

- **Resource** – You use an Amazon Resource Name (ARN) to identify the resource that the policy applies to. CloudWatch does not have any resources for you to control using policies resources, so you always use the wildcard character (*) in IAM policies. For more information, see [CloudWatch Resources and Operations \(p. 171\)](#).
- **Action** – You use action keywords to identify resource operations that you want to allow or deny. For example, the `cloudwatch:ListMetrics` permission allows the user permissions to perform the `ListMetrics` operation.
- **Effect** – You specify the effect, either allow or deny, when the user requests the specific action. If you don't explicitly grant access to (allow) a resource, access is implicitly denied. You can also explicitly deny access to a resource, which you might do to make sure that a user cannot access it, even if a different policy grants access.
- **Principal** – In identity-based policies (IAM policies), the user that the policy is attached to is the implicit principal. For resource-based policies, you specify the user, account, service, or other entity that you want to receive permissions (applies to resource-based policies only). CloudWatch doesn't support resource-based policies.

To learn more about IAM policy syntax and descriptions, see [AWS IAM Policy Reference](#) in the *IAM User Guide*.

For a table showing all of the CloudWatch API actions and the resources that they apply to, see [Amazon CloudWatch Permissions Reference \(p. 180\)](#).

Specifying Conditions in a Policy

When you grant permissions, you can use the access policy language to specify the conditions when a policy should take effect. For example, you might want a policy to be applied only after a specific date. For more information about specifying conditions in a policy language, see [Condition](#) in the *IAM User Guide*.

To express conditions, you use predefined condition keys. There are no condition keys specific to CloudWatch. However, there are AWS-wide condition keys that you can use as appropriate. For a complete list of AWS-wide keys, see [Available Keys for Conditions](#) in the *IAM User Guide*.

Using Identity-Based Policies (IAM Policies) for CloudWatch

This topic provides examples of identity-based policies that demonstrate how an account administrator can attach permissions policies to IAM identities (that is, users, groups, and roles) and thereby grant permissions to perform operations on CloudWatch resources.

Important

We recommend that you first review the introductory topics that explain the basic concepts and options available to manage access to your CloudWatch resources. For more information, see [Access Control](#) (p. 170).

The sections in this topic cover the following:

- [Permissions Required to Use the CloudWatch Console](#) (p. 175)
- [AWS Managed \(Predefined\) Policies for CloudWatch](#) (p. 178)
- [Customer Managed Policy Examples](#) (p. 178)

The following shows an example of a permissions policy.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [ "cloudwatch:GetMetricStatistics", "cloudwatch:ListMetrics" ],
    "Resource": "*",
    "Condition": {
      "Bool": {
        "aws:SecureTransport": "true"
      }
    }
  }
]
```

This sample policy has one statement that grants permissions to a group for two CloudWatch actions (`cloudwatch:GetMetricStatistics` and `cloudwatch:ListMetrics`), but only if the group uses SSL with the request (`"aws:SecureTransport": "true"`). For more information about the elements within an IAM policy statement, see [Specifying Policy Elements: Actions, Effects, and Principals](#) (p. 174) and [IAM Policy Elements Reference](#) in *IAM User Guide*.

Permissions Required to Use the CloudWatch Console

For a user to work with the CloudWatch console, that user must have a minimum set of permissions that allows the user to describe other AWS resources in their AWS account. The CloudWatch console requires permissions from the following services:

- Auto Scaling
- CloudTrail
- CloudWatch
- CloudWatch Events

- CloudWatch Logs
- Amazon EC2
- Amazon ES
- IAM
- Amazon Kinesis
- Lambda
- Amazon S3
- Amazon SNS
- Amazon SQS
- Amazon SWF

If you create an IAM policy that is more restrictive than the minimum required permissions, the console won't function as intended for users with that IAM policy. To ensure that those users can still use the CloudWatch console, also attach the `CloudWatchReadOnlyAccess` managed policy to the user, as described in [AWS Managed \(Predefined\) Policies for CloudWatch \(p. 178\)](#).

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the CloudWatch API.

The full set of permissions required to work with the CloudWatch console are listed below:

- `applicationautoscaling:describeScalingPolicies`
- `autoscaling:describeAutoScalingGroups`
- `autoscaling:describePolicies`
- `cloudtrail:describeTrails`
- `cloudwatch:deleteAlarms`
- `cloudwatch:describeAlarmHistory`
- `cloudwatch:describeAlarms`
- `cloudwatch:getMetricData`
- `cloudwatch:getMetricDataForAccounts`
- `cloudwatch:getMetricStatistics`
- `cloudwatch:listMetrics`
- `cloudwatch:putMetricAlarm`
- `cloudwatch:putMetricData`
- `ec2:describeInstances`
- `ec2:describeTags`
- `ec2:describeVolumes`
- `es:describeElasticsearchDomain`
- `es:listDomainNames`
- `events:deleteRule`
- `events:describeRule`
- `events:disableRule`
- `events:enableRule`
- `events:listRules`
- `events:putRule`
- `iam:attachRolePolicy`
- `iam:createRole`
- `iam:getPolicy`
- `iam:getPolicyVersion`

- iam:getRole
- iam:listAttachedRolePolicies
- iam:listRoles
- kinesis:describeStreams
- kinesis:listStreams
- lambda:addPermission
- lambda:createFunction
- lambda:getFunctionConfiguration
- lambda:listAliases
- lambda:listFunctions
- lambda:listVersionsByFunction
- lambda:removePermission
- logs:cancelExportTask
- logs:createExportTask
- logs:createLogGroup
- logs:createLogStream
- logs:deleteLogGroup
- logs:deleteLogStream
- logs:deleteMetricFilter
- logs:deleteRetentionPolicy
- logs:deleteSubscriptionFilter
- logs:describeExportTasks
- logs:describeLogGroups
- logs:describeLogStreams
- logs:describeMetricFilters
- logs:describeSubscriptionFilters
- logs:filterLogEvents
- logs:getLogEvents
- logs:putMetricFilter
- logs:putRetentionPolicy
- logs:putSubscriptionFilter
- logs:testMetricFilter
- s3:createBucket
- s3:listBuckets
- sns:createTopic
- sns:getTopicAttributes
- sns:listSubscriptions
- sns:listTopics
- sns:setTopicAttributes
- sns:subscribe
- sns:unsubscribe
- sqs:getQueueAttributes
- sqs:getQueueUrl
- sqs:listQueues
- sqs:setQueueAttributes
- swf:createAction

- `swf:describeAction`
- `swf:listActionTemplates`
- `swf:registerAction`
- `swf:registerDomain`
- `swf:updateAction`

AWS Managed (Predefined) Policies for CloudWatch

AWS addresses many common use cases by providing standalone IAM policies that are created and administered by AWS. These AWS managed policies grant necessary permissions for common use cases so that you can avoid having to investigate what permissions are needed. For more information, see [AWS Managed Policies](#) in the *IAM User Guide*.

The following AWS managed policies, which you can attach to users in your account, are specific to CloudWatch:

- **CloudWatchFullAccess** – Grants full access to CloudWatch.
- **CloudWatchReadOnlyAccess** – Grants read-only access to CloudWatch.
- **CloudWatchActionsEC2Access** – Grants read-only access to CloudWatch alarms and metrics as well as Amazon EC2 metadata. Grants access to the Stop, Terminate, and Reboot API actions for EC2 instances.

Note

You can review these permissions policies by signing in to the IAM console and searching for specific policies there.

You can also create your own custom IAM policies to allow permissions for CloudWatch actions and resources. You can attach these custom policies to the IAM users or groups that require those permissions.

Customer Managed Policy Examples

In this section, you can find example user policies that grant permissions for various CloudWatch actions. These policies work when you are using the CloudWatch API, AWS SDKs, or the AWS CLI.

Examples

- [Example 1: Allow User Full Access to CloudWatch \(p. 178\)](#)
- [Example 2: Allow Read-Only Access to CloudWatch \(p. 179\)](#)
- [Example 3: Stop or Terminate an Amazon EC2 Instance \(p. 179\)](#)

Example 1: Allow User Full Access to CloudWatch

The following policy allows a user to access all CloudWatch actions, CloudWatch Logs actions, Amazon SNS actions, and read-only access to Auto Scaling.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
```

```
        "autoscaling:Describe*",
        "cloudwatch:*",
        "logs:*",
        "sns:*"
    ],
    "Effect": "Allow",
    "Resource": "*"
}
]
```

Example 2: Allow Read-Only Access to CloudWatch

The following policy allows a user read-only access to CloudWatch and view Auto Scaling actions, CloudWatch metrics, CloudWatch Logs data, and alarm-related Amazon SNS data.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "autoscaling:Describe*",
        "cloudwatch:Describe*",
        "cloudwatch:Get*",
        "cloudwatch:List*",
        "logs:Get*",
        "logs:Describe*",
        "sns:Get*",
        "sns:List*"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

Example 3: Stop or Terminate an Amazon EC2 Instance

The following policy allows an CloudWatch alarm action to stop or terminate an EC2 instance. In the sample below, the GetMetricStatistics, ListMetrics, and DescribeAlarms actions are optional. It is recommended that you include these actions to ensure that you have correctly stopped or terminated the instance.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "cloudwatch:PutMetricAlarm",
        "cloudwatch:GetMetricStatistics",
        "cloudwatch:ListMetrics",
        "cloudwatch:DescribeAlarms"
      ],
      "Sid": "0000000000000000",
      "Resource": [
        "*"
      ]
    }
  ]
}
```

```

    ],
    "Effect": "Allow"
  },
  {
    "Action": [
      "ec2:DescribeInstanceStatus",
      "ec2:DescribeInstances",
      "ec2:StopInstances",
      "ec2:TerminateInstances"
    ],
    "Sid": "0000000000000000",
    "Resource": [
      "*"
    ],
    "Effect": "Allow"
  }
]
}

```

Amazon CloudWatch Permissions Reference

When you are setting up [Access Control \(p. 170\)](#) and writing permissions policies that you can attach to an IAM identity (identity-based policies), you can use the following table as a reference. The table lists each CloudWatch API operation and the corresponding actions for which you can grant permissions to perform the action. You specify the actions in the policy's `Action` field, and you specify a wildcard character (*) as the resource value in the policy's `Resource` field.

You can use AWS-wide condition keys in your CloudWatch policies to express conditions. For a complete list of AWS-wide keys, see [Available Keys](#) in the *IAM User Guide*.

Note

To specify an action, use the `cloudwatch:` prefix followed by the API operation name. For example: `cloudwatch:GetMetricStatistics`, `cloudwatch:ListMetrics`, or `cloudwatch:*` (for all CloudWatch actions).

Tables

- [CloudWatch API Operations and Required Permissions](#)
- [CloudWatch Events API Operations and Required Permissions](#)
- [CloudWatch Logs API Operations and Required Permissions](#)
- [Amazon EC2 API Operations and Required Permissions](#)
- [Auto Scaling API Operations and Required Permissions](#)

CloudWatch API Operations and Required Permissions for Actions

CloudWatch API Operations	Required Permissions (API Actions)
DeleteAlarms	<code>cloudwatch:DeleteAlarms</code> Required to delete an alarm.
DescribeAlarmHistory	<code>cloudwatch:DescribeAlarmHistory</code> Required to view alarm history.
DescribeAlarms	<code>cloudwatch:DescribeAlarms</code>

CloudWatch API Operations	Required Permissions (API Actions)
	Required to retrieve alarm information by name.
DescribeAlarmsForMetric	<i>cloudwatch:DescribeAlarmsForMetric</i> Required to view alarms for a metric.
DisableAlarmActions	<i>cloudwatch:DisableAlarmActions</i> Required to disable an alarm action.
EnableAlarmActions	<i>cloudwatch:EnableAlarmActions</i> Required to enable an alarm action.
GetMetricData	<i>cloudwatch:GetMetricData</i> Required to view or list dashboards and view metric data in dashboard widgets.
GetMetricStatistics	<i>cloudwatch:GetMetricStatistics</i> Required to view graphs in other parts of the CloudWatch console and in dashboard widgets.
ListMetrics	<i>cloudwatch:ListMetrics</i> Required to view or search metric names within the CloudWatch console and in the CLI. Required to select metrics on dashboard widgets.
PutMetricAlarm	<i>cloudwatch:PutMetricAlarm</i> Required to create or update an alarm.
PutMetricData	<i>cloudwatch:PutMetricData</i> Required to create metrics and create or delete dashboards.
SetAlarmState	<i>cloudwatch:SetAlarmState</i> Required to manually set an alarm's state.

CloudWatch Events API Operations and Required Permissions for Actions

CloudWatch Events API Operations	Required Permissions (API Actions)
DeleteRule	<i>events:DeleteRule</i> Required to delete a rule.
DescribeRule	<i>events:DescribeRule</i> Required to list the details about a rule.
DisableRule	<i>events:DisableRule</i> Required to disable a rule.
EnableRule	<i>events:EnableRule</i>

CloudWatch Events API Operations	Required Permissions (API Actions)
	Required to enable a rule.
ListRuleNamesByTarget	<i>events:ListRuleNamesByTarget</i> Required to list rules associated with a target.
ListRules	<i>events:ListRules</i> Required to list all rules in your account.
ListTargetsByRule	<i>events:ListTargetsByRule</i> Required to list all targets associated with a rule.
PutEvents	<i>events:PutEvents</i> Required to add custom events that can be matched to rules.
PutRule	<i>events:PutRule</i> Required to create or update a rule.
PutTargets	<i>events:PutTargets</i> Required to add targets to a rule.
RemoveTargets	<i>events:RemoveTargets</i> Required to remove a target from a rule.
TestEventPattern	<i>events:TestEventPattern</i> Required to test an event pattern against a given event.

CloudWatch Logs API Operations and Required Permissions for Actions

CloudWatch Logs API Operations	Required Permissions (API Actions)
CancelExportTask	<i>logs:CancelExportTask</i> Required to cancel a pending or running export task.
CreateExportTask	<i>logs:CreateExportTask</i> Required to export data from a log group to an Amazon S3 bucket.
CreateLogGroup	<i>logs:CreateLogGroup</i> Required to create a new log group.
CreateLogStream	<i>logs:CreateLogStream</i> Required to create a new log stream in a log group.
DeleteDestination	<i>logs>DeleteDestination</i>

CloudWatch Logs API Operations	Required Permissions (API Actions)
	Required to delete a log destination and disables any subscription filters to it.
DeleteLogGroup	<i>logs:DeleteLogGroup</i> Required to delete a log group and any associated archived log events.
DeleteLogStream	<i>logs:DeleteLogStream</i> Required to delete a log stream and any associated archived log events.
DeleteMetricFilter	<i>logs:DeleteMetricFilter</i> Required to delete a metric filter associated with a log group.
DeleteRetentionPolicy	<i>logs:DeleteRetentionPolicy</i> Required to delete a log group's retention policy.
DeleteSubscriptionFilter	<i>logs:DeleteSubscriptionFilter</i> Required to delete the subscription filter associated with a log group.
DescribeDestinations	<i>logs:DescribeDestinations</i> Required to view all destinations associated with the account.
DescribeExportTasks	<i>logs:DescribeExportTasks</i> Required to view all export tasks associated with the account.
DescribeLogGroups	<i>logs:DescribeLogGroups</i> Required to view all log groups associated with the account.
DescribeLogStreams	<i>logs:DescribeLogStreams</i> Required to view all log streams associated with a log group.
DescribeMetricFilters	<i>logs:DescribeMetricFilters</i> Required to view all metrics associated with a log group.
DescribeSubscriptionFilters	<i>logs:DescribeSubscriptionFilters</i> Required to view all subscription filters associated with a log group.

CloudWatch Logs API Operations	Required Permissions (API Actions)
FilterLogEvents	<i>logs:FilterLogEvents</i> Required to sort log events by log group filter pattern.
GetLogEvents	<i>logs:GetLogEvents</i> Required to retrieve log events from a log stream.
PutDestination	<i>logs:PutDestination</i> Required to create or update a destination log stream (such as an Amazon Kinesis stream).
PutDestinationPolicy	<i>logs:PutDestinationPolicy</i> Required to create or update an access policy associated with an existing log destination.
PutLogEvents	<i>logs:PutLogEvents</i> Required to upload a batch of log events to a log stream.
PutMetricFilter	<i>logs:PutMetricFilter</i> Required to create or update a metric filter and associate it with a log group.
PutRetentionPolicy	<i>logs:PutRetentionPolicy</i> Required to set the number of days to keep log events (retention) in a log group.
PutSubscriptionFilter	<i>logs:PutSubscriptionFilter</i> Required to create or update a subscription filter and associate it with a log group.
TestMetricFilter	<i>logs:TestMetricFilter</i> Required to test a filter pattern against a sampling of log event messages.

Amazon EC2 API Operations and Required Permissions for Actions

Amazon EC2 API Operations	Required Permissions (API Actions)
DescribeInstanceStatus	<i>ec2:DescribeInstanceStatus</i> Required to view EC2 instance status details.
DescribeInstances	<i>ec2:DescribeInstances</i> Required to view EC2 instance details.
RebootInstances	<i>ec2:RebootInstances</i>

Amazon EC2 API Operations	Required Permissions (API Actions)
	Required to reboot an EC2 instance.
StopInstances	<i>ec2:StopInstances</i> Required to stop an EC2 instance.
TerminateInstances	<i>ec2:TerminateInstances</i> Required to terminate an EC2 instance.

Auto Scaling API Operations and Required Permissions for Actions

Auto Scaling API Operations	Required Permissions (API Actions)
Scaling	<i>autoscaling:Scaling</i> Required to scale an Auto Scaling group.
Trigger	<i>autoscaling:Trigger</i> Required to trigger an Auto Scaling action.

Logging Amazon CloudWatch API Calls in AWS CloudTrail

AWS CloudTrail is a service that captures API calls made by or on behalf of your AWS account. This information is collected and written to log files that are stored in an Amazon S3 bucket that you specify. API calls are logged whenever you use the API, the console, or the AWS CLI. Using the information collected by CloudTrail, you can determine what request was made, the source IP address the request was made from, who made the request, when it was made, and so on.

To learn more about CloudTrail, including how to configure and enable it, see the [What is AWS CloudTrail](#) in the *AWS CloudTrail User Guide*.

Topics

- [CloudWatch Information in CloudTrail](#) (p. 186)
- [Understanding Log File Entries](#) (p. 188)

CloudWatch Information in CloudTrail

If CloudTrail logging is turned on, calls made to API actions are captured in log files. Every log file entry contains information about who generated the request. For example, if a request is made to create or update a CloudWatch alarm (`PutMetricAlarm`), CloudTrail logs the user identity of the person or service that made the request.

The user identity information in the log entry helps you determine the following:

- Whether the request was made with root or IAM user credentials
- Whether the request was made with temporary security credentials for a role or federated user
- Whether the request was made by another AWS service

For more information, see the [CloudTrail userIdentity Element](#) in the *AWS CloudTrail User Guide*.

You can store your log files in your bucket for as long as you want, but you can also define Amazon S3 lifecycle rules to archive or delete log files automatically. By default, your log files are encrypted by using Amazon S3 server-side encryption (SSE).

If you want to be notified upon log file delivery, you can configure CloudTrail to publish Amazon SNS notifications when new log files are delivered. For more information, see [Configuring Amazon SNS Notifications for CloudTrail](#) in the *AWS CloudTrail User Guide*.

You can also aggregate Amazon CloudWatch Logs log files from multiple AWS regions and multiple AWS accounts into a single Amazon S3 bucket. For more information, see [Receiving CloudTrail Log Files from Multiple Regions](#) and [Receiving CloudTrail Log Files from Multiple Accounts](#) in the *AWS CloudTrail User Guide*.

When logging is turned on, the following API actions are written to CloudTrail:

CloudWatch

- DeleteAlarms
- DescribeAlarmHistory
- DescribeAlarms
- DescribeAlarmsForMetric
- DisableAlarmActions
- EnableAlarmActions
- PutMetricAlarm
- SetAlarmState

The CloudWatch `GetMetricStatistics`, `ListMetrics`, and `PutMetricData` API actions are not supported. For more information about all of these actions, see the [Amazon CloudWatch API Reference](#).

CloudWatch Events

- DeleteRule
- DescribeRule
- DisableRule
- EnableRule
- ListRuleNamesByTarget
- ListRules
- ListTargetsByRule
- PutRule
- PutTargets
- RemoveTargets
- TestEventPattern

For more information about these actions, see the [Amazon CloudWatch Events API Reference](#).

CloudWatch Logs

Request and response elements are logged for these API actions:

- CancelExportTask
- CreateExportTask
- CreateLogGroup
- CreateLogStream
- DeleteDestination
- DeleteLogGroup

- DeleteLogStream
- DeleteMetricFilter
- DeleteRetentionPolicy
- DeleteSubscriptionFilter
- PutDestination
- PutDestinationPolicy
- PutMetricFilter
- PutRetentionPolicy
- PutSubscriptionFilter
- TestMetricFilter

Only Request elements are logged for these API actions:

- DescribeDestinations
- DescribeExportTasks
- DescribeLogGroups
- DescribeLogStreams
- DescribeMetricFilters
- DescribeSubscriptionFilters

The CloudWatch Logs `GetLogEvents`, `PutLogEvents`, and `FilterLogEvents` API actions are not supported. For more information about these actions, see the [Amazon CloudWatch Logs API Reference](#).

Understanding Log File Entries

CloudTrail log files contain one or more log entries. Each entry lists multiple JSON-formatted events. A log entry represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. The log entries are not an ordered stack trace of the public API calls, so they do not appear in any specific order. Log file entries for all API actions are similar to the examples below.

The following log file entry shows that a user called the CloudWatch **PutMetricAlarm** action.

```
{
  "Records": [{
    "eventVersion": "1.01",
    "userIdentity": {
      "type": "Root",
      "principalId": "EX_PRINCIPAL_ID",
      "arn": "arn:aws:iam::123456789012:root",
      "accountId": "123456789012",
      "accessKeyId": "EXAMPLE_KEY_ID"
    },
    "eventTime": "2014-03-23T21:50:34Z",
    "eventSource": "monitoring.amazonaws.com",
    "eventName": "PutMetricAlarm",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "127.0.0.1",
    "userAgent": "aws-sdk-ruby2/2.0.0.rc4 ruby/1.9.3 x86_64-linux
Seahorse/0.1.0",
```

```
    "requestParameters": {
      "threshold": 50.0,
      "period": 60,
      "metricName": "CloudTrail Test",
      "evaluationPeriods": 3,
      "comparisonOperator": "GreaterThanThreshold",
      "namespace": "AWS/CloudWatch",
      "alarmName": "CloudTrail Test Alarm",
      "statistic": "Sum"
    },
    "responseElements": null,
    "requestID": "29184022-b2d5-11e3-a63d-9b463e6d0ff0",
    "eventID": "b096d5b7-dcf2-4399-998b-5a53eca76a27"
  },
  ..additional entries
]
}
```

The following log file entry shows that a user called the CloudWatch Events **PutRule** action.

```
{
  "eventVersion": "1.03",
  "userIdentity": {
    "type": "Root",
    "principalId": "123456789012",
    "arn": "arn:aws:iam::123456789012:root",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2015-11-17T23:56:15Z"
      }
    }
  },
  "eventTime": "2015-11-18T00:11:28Z",
  "eventSource": "events.amazonaws.com",
  "eventName": "PutRule",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "AWS Internal",
  "userAgent": "AWS CloudWatch Console",
  "requestParameters": {
    "description": "",
    "name": "cttest2",
    "state": "ENABLED",
    "eventPattern": "{\"source\": [\"aws.ec2\"], \"detail-type\": [\"EC2 Instance State-change Notification\"]}",
    "scheduleExpression": ""
  },
  "responseElements": {
    "ruleArn": "arn:aws:events:us-east-1:123456789012:rule/cttest2"
  },
  "requestID": "e9caf887-8d88-11e5-a331-3332aa445952",
  "eventID": "49d14f36-6450-44a5-a501-b0fdcdfaeb98",
  "eventType": "AwsApiCall",
  "apiVersion": "2015-10-07",
  "recipientAccountId": "123456789012"
}
```

The following log file entry shows that a user called the CloudWatch Logs **CreateExportTask** action.

```
{
  "eventVersion": "1.03",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "EX_PRINCIPAL_ID",
    "arn": "arn:aws:iam::123456789012:user/someuser",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "userName": "someuser"
  },
  "eventTime": "2016-02-08T06:35:14Z",
  "eventSource": "logs.amazonaws.com",
  "eventName": "CreateExportTask",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "127.0.0.1",
  "userAgent": "aws-sdk-ruby2/2.0.0.rc4 ruby/1.9.3 x86_64-linux
Seahorse/0.1.0",
  "requestParameters": {
    "destination": "yourdestination",
    "logGroupName": "yourloggroup",
    "to": 123456789012,
    "from": 0,
    "taskName": "yourtask"
  },
  "responseElements": {
    "taskId": "15e5e534-9548-44ab-a221-64d9d2b27b9b"
  },
  "requestID": "1cd74c1c-ce2e-12e6-99a9-8dbb26bd06c9",
  "eventID": "fd072859-bd7c-4865-9e76-8e364e89307c",
  "eventType": "AwsApiCall",
  "apiVersion": "20140328",
  "recipientAccountId": "123456789012"
}
```

Document History

The following table describes the important changes to the Amazon CloudWatch User's Guide.

Change	Description	Release Date
Added metrics for Amazon Polly	Added metrics for Amazon Polly. For more information, see Amazon Polly Metrics (p. 118) .	1 December 2016
Added metrics for Amazon Kinesis Analytics	Added metrics for Amazon Kinesis Analytics. For more information, see Amazon Kinesis Analytics Metrics (p. 105) .	1 December 2016
Added support for percentile statistics	You can specify any percentile, using up to two decimal places (for example, p95.45). For more information, see Percentiles (p. 7) .	17 November 2016
Added metrics for Amazon Simple Email Service	Added metrics for Amazon Simple Email Service. For more information, see Amazon Simple Email Service Metrics and Dimensions (p. 126) .	2 November 2016
Updated metrics retention	Amazon CloudWatch now retains metrics data for 15 months instead of 14 days.	1 November 2016
Updated metrics console interface	The CloudWatch console is updated with improvements to existing functionality and new functionality.	1 November 2016
Added metrics for Amazon Elastic Transcoder	Added metrics for Amazon Elastic Transcoder. For more information, see Amazon Elastic Transcoder Metrics and Dimensions (p. 101) .	20 September 2016
Added metrics for Amazon API Gateway	Added metrics for Amazon API Gateway. For more information, see Amazon API Gateway Metrics and Dimensions (p. 48) .	9 September 2016
Added metrics for AWS Key Management Service	Added metrics for AWS Key Management Service. For more information, see AWS Key Management Service Metrics and Dimensions (p. 114) .	9 September 2016
Added metrics for the new Application	Added metrics for Application Load Balancers. For more information, see Elastic Load Balancing Metrics and Dimensions (p. 84) .	11 August 2016

Change	Description	Release Date
Load Balancers supported by Elastic Load Balancing		
Added new NetworkPacketsIn and NetworkPacketsOut metrics for Amazon EC2	Added new NetworkPacketsIn and NetworkPacketsOut metrics for Amazon EC2. For more information, see Amazon EC2 Metrics and Dimensions (p. 66).	23 March 2016
Added new metrics for Amazon EC2 Spot fleet	Added new metrics for Amazon EC2 Spot fleet. For more information, see Amazon EC2 Spot Fleet Metrics and Dimensions (p. 70).	21 March 2016
Added new CloudWatch Logs metrics	Added new CloudWatch Logs metrics. For more information, see Amazon CloudWatch Logs Metrics and Dimensions (p. 55).	10 March 2016
Added Amazon Elasticsearch Service and AWS WAF metrics and dimensions	Added Amazon Elasticsearch Service and AWS WAF metrics and dimensions. For more information, see Amazon Elasticsearch Service Metrics and Dimensions (p. 99) and AWS WAF Metrics and Dimensions (p. 145).	14 October 2015
Added support for CloudWatch dashboards	Dashboards are customizable home pages in the CloudWatch console that you can use to monitor your resources in a single view, even those that are spread out across different regions. For more information, see Using Amazon CloudWatch Dashboards (p. 18).	8 October 2015
Added AWS Lambda metrics and dimensions	Added AWS Lambda metrics and dimensions. For more information, see AWS Lambda Metrics and Dimensions (p. 114).	4 September 2015
Added Amazon EC2 Container Service metrics and dimensions	Added Amazon EC2 Container Service metrics and dimensions. For more information, see Amazon ECS Metrics and Dimensions (p. 71).	17 August 2015
Added Amazon Simple Storage Service metrics and dimensions	Added Amazon Simple Storage Service metrics and dimensions. For more information, see Amazon Simple Storage Service Metrics and Dimensions (p. 130).	26 July 2015
New feature: Reboot alarm action	Added the reboot alarm action and new IAM role for use with alarm actions. For more information, see Create Alarms to Stop, Terminate, Reboot, or Recover an Instance (p. 159).	23 July 2015
Added Amazon WorkSpaces metrics and dimensions	Added Amazon WorkSpaces metrics and dimensions. For more information, see Amazon WorkSpaces Metrics and Dimensions (p. 146).	30 April 2015

Change	Description	Release Date
Added Amazon Machine Learning metrics and dimensions	Added Amazon Machine Learning metrics and dimensions. For more information, see Amazon Machine Learning Metrics and Dimensions (p. 116) .	9 April 2015
New feature: Amazon EC2 instance recovery alarm actions	Updated alarm actions to include new EC2 instance recovery action. For more information, see Create Alarms to Stop, Terminate, Reboot, or Recover an Instance (p. 159) .	12 March 2015
Added Amazon CloudFront and Amazon CloudSearch metrics and dimensions	Added Amazon CloudFront and Amazon CloudSearch metrics and dimensions. For more information, see Amazon CloudFront Metrics and Dimensions (p. 51) and Amazon CloudSearch Metrics and Dimensions (p. 53) .	6 March 2015
Added Amazon Simple Workflow Service metrics and dimensions	Added Amazon Simple Workflow Service metrics and dimensions. For more information, see Amazon SWF Metrics and Dimensions (p. 133) .	9 May 2014
Updated guide to add support for AWS CloudTrail	Added a new topic to explain how you can use AWS CloudTrail to log activity in Amazon CloudWatch. For more information, see Logging Amazon CloudWatch API Calls in AWS CloudTrail (p. 186) .	30 April 2014
Updated guide to use the new AWS Command Line Interface (AWS CLI)	The AWS CLI is a cross-service CLI with a simplified installation, unified configuration, and consistent command line syntax. The AWS CLI is supported on Linux/Unix, Windows, and Mac. The CLI examples in this guide have been updated to use the new AWS CLI. For information about how to install and configure the new AWS CLI, see Getting Set Up with the AWS Command Line Interface in the <i>AWS Command Line Interface User Guide</i> .	21 February 2014
Added Amazon Redshift and AWS OpsWorks metrics and dimensions	Added Amazon Redshift and AWS OpsWorks metrics and dimensions. For more information, see Amazon Redshift Metrics and Dimensions (p. 119) and AWS OpsWorks Metrics and Dimensions (p. 117) .	16 July 2013
Added Amazon Route 53 metrics and dimensions	Added Amazon Route 53 metrics and dimensions. For more information, see Amazon Route 53 Metrics and Dimensions (p. 124) .	26 June 2013
New feature: Amazon CloudWatch Alarm Actions	Added a new section to document Amazon CloudWatch alarm actions, which you can use to stop or terminate an Amazon Elastic Compute Cloud instance. For more information, see Create Alarms to Stop, Terminate, Reboot, or Recover an Instance (p. 159) .	8 January 2013

Change	Description	Release Date
Updated EBS metrics	Updated the EBS metrics to include two new metrics for Provisioned IOPS volumes. For more information, see Amazon EBS Metrics and Dimensions (p. 80) .	20 November 2012
New billing alerts	You can now monitor your AWS charges using Amazon CloudWatch metrics and create alarms to notify you when you have exceeded the specified threshold. For more information, see Create a Billing Alarm to Monitor Your Estimated AWS Charges (p. 165) .	10 May 2012
New metrics	You can now access six new Elastic Load Balancing metrics that provide counts of various HTTP response codes. For more information, see Elastic Load Balancing Metrics and Dimensions (p. 84) .	19 October 2011
New feature	You can now access metrics from Amazon EMR. For more information, see Amazon EMR Metrics and Dimensions (p. 89) .	30 June 2011
New feature	You can now access metrics from Amazon Simple Notification Service and Amazon Simple Queue Service. For more information, see Amazon Simple Notification Service Metrics and Dimensions (p. 126) and Amazon SQS Metrics and Dimensions (p. 128) .	14 July 2011
New Feature	Added information about using the <code>PutMetricData</code> API to publish custom metrics. For more information, see Publish Custom Metrics (p. 43) .	10 May 2011
Updated metrics retention	Amazon CloudWatch now retains the history of an alarm for two weeks rather than six weeks. With this change, the retention period for alarms matches the retention period for metrics data.	07 April 2011
New feature	Added ability to send Amazon Simple Notification Service or Auto Scaling notifications when a metric has crossed a threshold. For more information, see Alarms (p. 7) .	02 December 2010
New feature	A number of CloudWatch actions now include the <code>MaxRecords</code> and <code>NextToken</code> parameters, which enable you to control pages of results to display.	02 December 2010
New feature	This service now integrates with AWS Identity and Access Management (IAM).	02 December 2010