

User Guide

Service Quotas



Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Service Quotas: User Guide

Copyright © 2024 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 Service Quotas?	1
Features of Service Quotas	1
Terminology in Service Quotas	2
Accessing Service Quotas	3
Getting started: Customize the Service Quotas dashboard	5
Viewing service quotas	6
Requesting a quota increase	11
Using the AWS Management Console to request an increase	11
Using the AWS CLI to request a quota increase	12
Verifying your quota request	18
Tagging resources	23
Supported resources	23
Tag restrictions	24
Enabling required permissions	24
Managing tags	25
Managing tags from the AWS Management Console	25
Managing tags (CLI or API)	26
Controlling access using tags	26
Using request templates	28
Security	30
Data protection	30
Encryption at rest	31
Encryption in transit	31
Logging and monitoring	32
Overview	
Logging Service Quotas APIs with CloudTrail	32
Using CloudWatch alarms	36
Identity and access management	37
Grant permissions using IAM policies	37
AWS managed policies	38
API actions for Service Quotas	44
Service Quotas resources	45
Resource-level permissions for Service Quotas	45
Condition keys for Service Quotas	45

	Predefined AWS managed policies for Service Quotas	46
	Compliance validation	46
	Resilience	47
	Infrastructure Security	48
Qι	otas for Service Quotas	49
Do	ocument history	53

What is Service Quotas?

With Service Quotas, you can view and manage your quotas for AWS services from a central location. Quotas, also referred to as limits in AWS services, are the maximum values for the resources, actions, and items in your AWS account. Each AWS service defines its quotas and establishes default values for those quotas. If your business needs aren't met by the default limit of service resources or operations that apply to an AWS account, resource, or an AWS Region, you might need to increase your service quota values. Service Quotas enables you to look up your service quotas and to request increases. AWS Support might approve, deny, or partially approve your requests.

Contents

- Features of Service Quotas
- Terminology in Service Quotas
- Accessing Service Quotas

Features of Service Quotas

Service Quotas provides the following features:

View your service quotas

The Service Quotas console provides quick access to the AWS default quota values for your account, across all AWS Regions. When you select a service in the Service Quotas console, you see the service's quotas and if that quota is adjustable at the AWS account level. *Applied quotas* are overrides, or increases for a specific quota, over the AWS default value.

Request a service quota increase

To see if a quota is adjustable, go into the console, navigate to AWS services, and select the service from the list. From the service's details page, view the **Adjustable** column.

Each adjustable quota says at which level the quota can be increased. For service quotas that are adjustable at the *account* level, you can use Service Quotas to request a quota increase.

You can also increase certain quotas at the *resource* level.

To request a quota increase in the Service Quotas console, select the service and the specific quota, and then choose **Request quota increase**. Increases do take some time to review,

Features of Service Quotas

process, and approve. You can also use Service Quotas API operations or the AWS CLI tools to request service quota increases.

View current utilization of resources

After your account becomes active for a period of time, you can view a graph of your resource utilization.

Terminology in Service Quotas

The following terms are important for understanding Service Quotas and how it works.

service quota

The maximum number of service resources or operations that apply to an AWS account or an AWS Region. The number of AWS Identity and Access Management (IAM) roles per account is an example of an account-based quota. The number of virtual private clouds (VPCs) per Region is an example of a Region-based quota. To determine whether a service quota is Region-specific, check the description of the service quota.

adjustable value

A quota value that can be increased.

applied quota

The updated quota value after a quota increase.

default value

The initial quota value established by AWS.

global quota

A service quota applied at an account level. Global quotas are available in all AWS Regions. You can request an increase to a global quota only from US East (N. Virginia) for Public AWS partition, AWS GovCloud (US-West) for AWS GovCloud (US) Regions, and China (Beijing) for AWS China Regions.

usage

The number of resources or operations in use for a service quota.

utilization

The percentage of a service quota in use. For example, if the quota value is 200 resources and 150 resources are in use, then the utilization is 75 percent.

quota context info

A structure that describes the context for a resource-level quota. For resource-level quotas, such as Instances per OpenSearch Service Domain, you can apply the quota value at the resource-level for each OpenSearch Service Domain in your AWS account. Together the attributes of this structure help you understand how the quota is implemented by AWS and how you can manage it.

context ID

Specifies the resource, or resources, to which the quota applies. The value for this field is either an Amazon Resource Name (ARN) or *. If the value is an ARN, the quota value applies to that resource. If the value is *, then the quota value applies to all resources of that specific type.

context scope

Specifies the scope to which the quota value is applied.

context scope type

Specifies the resource type to which the quota can be applied.

quota applied at level

Filters an API response to return applied quota values at either the account level, resource level, or all levels.

quota requested at level

Filters an API response to return quota requests at either the account level, resource level, or all levels.

Accessing Service Quotas

You can work with Service Quotas in the following ways:

AWS Management Console

<u>The Service Quotas console</u> is a browser-based interface that you can use to view and manage your service quotas. You can perform almost any task that's related to your service quotas

Accessing Service Quotas 3

by using the console. You can access Service Quotas from any AWS Management Console page by choosing it on the top navigation bar, or by searching for Service Quotas in the AWS Management Console.

AWS Command Line Interface tools

By using the AWS Command Line Interface tools, you can issue commands at your system's command line to perform Service Quotas and other AWS tasks. This can be a faster and more convenient approach than using the console. The command line tools also are useful if you want to build scripts that perform AWS tasks.

AWS provides two sets of command line tools: the <u>AWS Command Line Interface</u> and the <u>AWS Tools for Windows PowerShell</u>. For information about installing and using the AWS CLI, see the <u>AWS Command Line Interface User Guide</u>. For information about installing and using the Tools for Windows PowerShell, see the AWS Tools for Windows PowerShell User Guide.

You need AWS CLI version 2.13.20 or higher to view and manage resource-level quotas such as Instances per domain for Amazon OpenSearch Service.

AWS SDKs

The AWS SDKs consist of libraries and sample code for various programming languages and platforms (for example, <u>Java</u>, <u>Python</u>, <u>Ruby</u>, <u>.NET</u>, <u>iOS and Android</u>, and <u>others</u>). The SDKs include tasks such as cryptographically signing requests, managing errors, and retrying requests automatically. For more information about the AWS SDKs, including how to download and install them, see <u>Tools for Amazon Web Services</u>.

Accessing Service Quotas 4

Customizing the Service Quotas dashboard

When you open the Service Quotas console, the dashboard displays cards for up to nine AWS services. Each card lists the number of service quotas for that AWS service. Choosing a card opens a page displaying the quotas for the service. You can modify which services appear on the dashboard.

To modify the dashboard service cards

- 1. Sign in to the AWS Management Console and open the Service Quotas console at https://console.aws.amazon.com/servicequotas/home.
- 2. On the dashboard, choose Modify dashboard cards.
- 3. The services that are currently selected appear on the right. If you have selected nine services, you must remove a service before you can add a different service. For each service that you don't want on the dashboard, choose **Remove**.
- 4. To add a service to the dashboard, select it from **Choose services**.
- 5. When you have finished adding and removing services, choose **Save**.

Next steps

- View the AWS default value and applied values of a particular quota.
- For adjustable quotas, you can request a quota increase.

Viewing service quotas

Service Quotas enables you to review the AWS default value and applied values of a particular *quota*. Certain resource-level quotas, such as Instances per domain for Amazon OpenSearch Service, also display applied quota values per resource. Refer to <u>Terminology in Service Quotas</u> for detailed definitions of these values.

Your account's actual quota value may be less than the AWS default quota value if you recently created the account, or if you use the account minimally.

AWS Management Console

To view the quotas for a service

- 1. Sign in to the AWS Management Console and open the Service Quotas console at https://console.aws.amazon.com/servicequotas/home.
- 2. In the navigation pane, choose **AWS services**.
- 3. Select an AWS service from the list, or enter the name of the service in the search field.

The console displays details for the selected service's quota, including the **quota name**, **applied quota value**, **AWS default quota**, **utilization**, and if the quota is **adjustable** at the account-level or resource-level.

If the applied quota value or utilization is not available, the console displays **Not available**. You can request your applied quota value through the Support Center Console.

 Choose a specific quota name to view the Details page, which displays that quota's description, quota code, quota ARN, utilization, applied quota value, and the AWS default quota.

If applicable, the Details page also displays any **resource-level quota**, **alarms**, **quota increase request history**, and any of the quota's **tags**.

AWS CLI

Viewing default quota values

View the default values for the quotas for a specific AWS service.

Call the <u>ListDefaultServiceQuotas</u> operation with a service code. If you don't have the service code, you can get a list of supported services with the <u>ListServices</u> operation. The response includes the ServiceCode and ServiceName for each service. The ServiceCode for Amazon OpenSearch Service is es. The following CLI example retrieves default values for Amazon OpenSearch Service quotas.

```
$ aws service-quotas list-aws-default-service-quotas \
        --service-code es \
{
    "Quotas": [
        {
            "QuotaName": "Domains per Region",
            "Adjustable": true,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-076D529E",
            "Value": 100.0,
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-076D529E",
            "Unit": "None",
        },
        {
            "QuotaName": "Dedicated master instances per domain",
            "Adjustable": false,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/L-
AE676A72",
            "Value": 5.0,
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-AE676A72",
            "Unit": "None",
        },
            "QuotaName": "Warm instances per domain",
            "Adjustable": false,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-1F053E6F",
            "Value": 150.0,
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
```

```
"QuotaCode": "L-1F053E6F",
            "Unit": "None",
        },
        {
            "QuotaName": "Instances per domain",
            "Adjustable": true,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-6408ABDE",
            "Value": 80.0,
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-6408ABDE",
            "Unit": "None",
            "QuotaContext": {
                               "ContextScope": "RESOURCE",
                               "ContextScopeType":
 "AWS::OpenSearchService::Domain",
        }
    ]
}
```

Viewing applied quota values

View the applied quota values for a specified AWS service. For some quotas, only the default values are available. If the applied quota value isn't available for a quota, the quota is not returned in the response. If this happens, contact AWS Support for the applied quota value.

Call the <u>ListServiceQuotas</u> operation with a service code. You can choose to retrieve
all applied quota values either at the account-level, resource-level, or all levels by
passing ACCOUNT, RESOURCE, or ALL respectively as the value for the parameter
QuotaAppliedAtLevel.

The following CLI example retrieves all quota values applied at the account-level for OpenSearch Service.

```
"Quotas": [
        {
            "QuotaName": "Domains per Region",
            "Adjustable": true,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-076D529E",
            "Value": 100.0,
            "QuotaAppliedAtLevel": "ACCOUNT",
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-076D529E",
            "Unit": "None",
        },
        {
            "QuotaName": "Dedicated master instances per domain",
            "Adjustable": false,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/L-
AE676A72",
            "Value": 5.0,
            "QuotaAppliedAtLevel": "ACCOUNT",
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-AE676A72",
            "Unit": "None",
        },
            "QuotaName": "Warm instances per domain",
            "Adjustable": false,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-1F053E6F",
            "Value": 150.0,
            "QuotaAppliedAtLevel": "ACCOUNT",
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-1F053E6F",
            "Unit": "None",
        },
            "QuotaName": "Instances per domain",
            "Adjustable": true,
```

```
"QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-6408ABDE",
            "Value": 80.0,
            "QuotaAppliedAtLevel": "ACCOUNT",
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-6408ABDE",
            "Unit": "None",
            "QuotaContext": {
                              "ContextScope": "RESOURCE",
                              "ContextScopeType":
 "AWS::OpenSearchService::Domain",
                              "ContextId": "*"
             }
        }
    ]
}
```

Requesting a quota increase

For adjustable quotas, you can request a quota increase at the *account-level* or the *resource-level*. Smaller increases are usually automatically approved while larger requests are submitted to AWS Support. Larger increase requests take time to review, process, approve, and deploy. You can track your request case in the AWS Support console.



Note

Quota increase requests don't receive priority support. AWS Support can approve, deny, or partially approve your requests. If you have an urgent quota request, or if your quota increase request is denied, contact AWS Support for assistance.

Using the AWS Management Console to request an increase

Increase your quotas at the account or resource level in the Getting Started with the AWS Management Console.

To request a service quota increase

- Sign in to the AWS Management Console and open the Service Quotas console at https:// console.aws.amazon.com/servicequotas/home.
- In the navigation pane, choose AWS services. 2.
- 3. Choose an AWS service from the list, or enter the name of the service in the search box.
- If the quota is adjustable, you can request a quota increase at either the account-level or 4. resource-level based on the value listed in the **Adjustability** column.
 - Account-level Request a quota increase at the account-level for an account-level quota such as Domains per Region for Amazon OpenSearch Service. To do so, select the quota from the list and choose **Request increase at account-level**.
 - Resource-level Request a quota increase for a specific resource for a resource-level quota such as Instances per domain for Amazon OpenSearch Service. To do so, choose the quota name to view additional information about the quota. Under the Resource-level quotas section, select the resource for which you want to increase the quota value, and choose the **Request increase at resource-level** button.

5. For **Increase quota value**, enter the new value. The new value must be greater than the current value.

- 6. Choose **Request**.
- 7. To view any pending or recently resolved requests in the console, navigate to the **Request history** tab from the service's details page, or choose **Dashboard** from the navigation pane.

 For pending requests, choose the status of the request to open the request receipt. The initial status of a request is **Pending**. After the status changes to **Quota requested**, you'll see the case number with AWS Support. Choose the case number to open the ticket for your request.

Using the AWS CLI to request a quota increase

Requesting a quota increase using the AWS CLI requires you to provide Service Quotas with the necessary permission to create a support case on your behalf. You can provide this permission by attaching the <u>AWS managed policy</u> ServiceQuotasFullAccess to your IAM principal or adding iam:CreateServiceLinkedRole to your existing IAM policy.

Account-level increase request AWS CLI

To request a quota increase at the account-level

The RequestServiceQuotaIncrease operation, which submits the request, requires the quota code for the quota. So begin by getting the quota code.

The following example commands show how to request a quota increase at the account-level for the Amazon OpenSearch Service.

- Get the list of services supported by Service Quotas with the <u>ListServices</u> operation.
 The response includes the ServiceCode and ServiceName for each service. The ServiceCode for Amazon OpenSearch Service is es.
- 2. Get the list of Amazon OpenSearch Service quotas and their corresponding applied quota values at the account-level by calling the ListServiceQuotas operation with request parameters ServiceCode as es, and QuotaAppliedAtLevel as ACCOUNT. The response includes the QuotaName, QuotaCode, Value, and QuotaAppliedAtLevel for each quota of the Amazon OpenSearch Service that is applied at the account-level. If the value in the QuotaAppliedAtLevel field is ACCOUNT, then the Value represents the Applied quota value at the account-level. The following CLI example retrieves the quota code for a OpenSearch Service quota.

```
$ aws service-quotas list-service-quotas \
        --service-code es \
        --quota-applied-at-level ACCOUNT
{
    "Quotas": [
        {
            "QuotaName": "Domains per Region",
            "Adjustable": true,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-076D529E",
            "Value": 100.0,
            "QuotaAppliedAtLevel": "ACCOUNT",
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-076D529E",
            "Unit": "None",
        },
            "QuotaName": "Dedicated master instances per domain",
            "Adjustable": false,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/L-
AE676A72",
            "Value": 5.0,
            "QuotaAppliedAtLevel": "ACCOUNT",
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-AE676A72",
            "Unit": "None",
        },
            "QuotaName": "Warm instances per domain",
            "Adjustable": false,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-1F053E6F",
            "Value": 150.0,
            "QuotaAppliedAtLevel": "ACCOUNT",
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-1F053E6F",
            "Unit": "None",
```

```
},
        {
            "QuotaName": "Instances per domain",
            "Adjustable": true,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-6408ABDE",
            "Value": 80.0,
            "QuotaAppliedAtLevel": "ACCOUNT",
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-6408ABDE",
            "Unit": "None",
            "QuotaContext": {
                              "ContextScope": "RESOURCE",
                              "ContextScopeType":
 "AWS::OpenSearchService::Domain",
                              "ContextId": "*"
        }
    ]
}
```

3. Next, call the <u>RequestServiceQuotaIncrease</u> operation and specify the QuotaCode in the request parameter.

The following example requests an increase at the account-level in the Instances per domain quota to 100. It uses the required quota code, L-6408ABDE, to identify the quota. If the command completes successfully, the Status field in the response displays the current status of the request. The QuotaRequestedAtLevel field in the response specifies that this request applies to the account-level.

Note

You can't request a quota increase at the account-level for a resource-level quota through the AWS CLI. This operation results in the creation of a support case where you can provide the ARN to specify the resource on which the new quota value should apply. However, the Instances per domain quota for Amazon OpenSearch Service is an exception.

```
$ aws service-quotas request-service-quota-increase \
```

```
--service-code es \
        --quota-code L-6408ABDE \
        --desired-value 100
{
    "RequestedQuota": {
        "QuotaName": "Instances per domain",
        "Status": "PENDING",
        "DesiredValue": 100,
        "Created": 1580446904.067,
        "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-6408ABDE",
        "ServiceName": "Amazon OpenSearch Service",
        "GlobalQuota": false,
        "ServiceCode": "es",
        "QuotaCode": "L-6408ABDE",
        "Requester": "{\"accountId\":\"123456789012\",\"callerArn\":
\"arn:aws:iam::123456789012:root\"}",
        "QuotaRequestedAtLevel": "ACCOUNT"
        "Id": "a12345",
        "Unit": "None"
        "QuotaContext": {
                          "ContextId": "*"
                         "ContextScopeType": "AWS::OpenSearchService::Domain",
                         "ContextScope": "RESOURCE",
        }
    }
}
```

 To get the updated status of the request, use the <u>GetRequestedServiceQuotaChange</u>, <u>ListRequestedServiceQuotaChangeHistory</u> or <u>ListRequestedServiceQuotaChangeHistoryByQuota operations</u>.

Resource-level quota increase request AWS CLI

To request a quota increase at the resource-level

The RequestServiceQuotaIncrease operation, which submits the request, requires the quota code for the quota. So begin by getting the quota code. To request a quota increase for a specific resource, use the Amazon Resource Name (ARN) ResourceARN as the value for the ContextId parameter when you make your request.

The following example commands show how to request a resource-level quota increase for the OpenSearch Service.

- Get the list of services supported by Service Quotas with the <u>ListServices</u> operation.
 The response includes the ServiceCode and ServiceName for each service. The ServiceCode for Amazon OpenSearch Service is es.
- 2. Get the list of Amazon OpenSearch Service quotas and their corresponding applied quota values at the resource-level by calling the ListServiceQuotas operation with request parameters ServiceCode as es, and QuotaAppliedAtLevel as RESOURCE. The response includes the QuotaName, QuotaCode, Value, and QuotaAppliedAtLevel for each quota of the Amazon OpenSearch Service that is applied at the resource-level. If the value in the QuotaAppliedAtLevel field is RESOURCE, then the Value represents the Applied quota value at the resource-level. In this case, the response for this quota will also contain the QuotaContext structure which further specifies the ContextId or the ARN to which the quota value is applied. The following CLI example retrieves the quota code for a OpenSearch Service quota.

```
$ aws service-quotas list-service-quotas \
        --service-code es \
        --quota-applied-at-level RESOURCE
{
    "Ouotas": [
         {
            "QuotaName": "Instances per domain",
            "Adjustable": true,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-6408ABDE",
            "Value": 100.0,
            "QuotaAppliedAtLevel": "RESOURCE",
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-6408ABDE",
            "Unit": "None",
            "QuotaContext": {
                              "ContextScope": "RESOURCE",
                              "ContextScopeType":
 "AWS::OpenSearchService::Domain",
                              "ContextId": "arn:aws:esus-
east-1:123456789012:domain/opensearch-domain-1",
```

```
},
       {
       "QuotaName": "Instances per domain",
            "Adjustable": true,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-6408ABDE",
            "Value": 100.0,
            "QuotaAppliedAtLevel": "RESOURCE",
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-6408ABDE",
            "Unit": "None",
            "QuotaContext": {
                              "ContextScope": "RESOURCE",
                              "ContextScopeType":
 "AWS::OpenSearchService::Domain",
                              "ContextId": "arn:aws:esus-
east-1:123456789012:domain/opensearch-domain-2",
                    }
         }
    ]
}
```

 Next, call the <u>RequestServiceQuotaIncrease</u> operation and specify the ServiceCode, QuotaCode, ContextId, and DesiredValue request parameters.

The following example requests an increase in the Instances per domain quota to 100 for a specific Amazon OpenSearch Service domain with the ARN as arn:aws:es:us-east-1:123456789012:domain/opensearch-domain-1. If the command completes successfully, the Status field in the response displays the current status of the request. QuotaRequestedAtLevel field in the response contains the value RESOURCE which specifies that this request is for a specific resource.

```
"QuotaName": "Instances per domain",
        "Status": "PENDING",
        "DesiredValue": 200.0,
        "Created": 1580446904.067,
        "QuotaArn": "arn:aws:servicequotas:us-east-1:123456789012:es/
L-6408ABDE",
        "ServiceName": "Amazon OpenSearch Service",
        "GlobalQuota": false,
        "ServiceCode": "es",
        "QuotaCode": "L-6408ABDE",
        "Requester": "{\"accountId\":\"123456789012\",\"callerArn\":
\"arn:aws:iam::123456789012:root\"}",
        "QuotaRequestedAtLevel": "RESOURCE",
        "Id": "a12345",
        "Unit": "None"
        "QuotaContext": {
                          "ContextId": "arn:aws:es:us-east-1:123456789012:domain/
opensearch-domain-1"
                         "ContextScopeType": "AWS::OpenSearchService::Domain",
                         "ContextScope": "RESOURCE",
             }
    }
}
```

To get the updated status of the request, use the GetRequestedServiceQuotaChange, ListRequestedServiceQuotaChangeHistory or $\underline{List Requested Service Quot} a Change History By Quota\ operations.$

After the request is resolved, the **Applied quota value** for the quota is set to the new value.

Verifying your quota request

Verify your quota request by viewing the request history history in the Service Quotas console. The console displays all open quota requests and requests closed in the last 90 days.



Note

Some AWS services may only be available in certain AWS Regions. If you have quota increase requests in different Regions, be sure to select the appropriate Region first.


Using the AWS Management Console

To view the quota request history

1. Sign in to the AWS Management Console and open the Service Quotas console at https://console.aws.amazon.com/servicequotas/home.

2. To view any pending or recently resolved requests, choose **Quota request history** from the navigation pane.

The **Recent quota increase requests** panel displays information about your open recent quota increase requests and any requests closed within 90 days.



- **Service** Displays the service name selected for the request.
- Quota name Displays the quota name selected for the quota increase.
- Status Displays the status of a request for a quota increase.

You may see the following status types:

- Requested AWS received your quota increase request.
- Pending Quota increase request is under review by AWS.
- **Case_Opened** Service Quotas has opened a support case to process the request. Please follow-up on the support case for more information.
- Approved Quota increase request is approved.
- Denied Quota increase request can't be approved by Service Quotas. Please contact AWS Support for more details.
- Not_Approved Quota increase request can't be approved by Service Quotas. Please contact AWS Support for more details.
- **Case_Closed** The support case associated with this request was closed. Check the support case correspondence for more information.

• Invalid_Request – Service Quotas can't process your resource-level quota increase request because the ResourceARN specified as part of the ContextId attribute is invalid.

- Requested quota value The increased quota value you requested for the quota.
- **Request date** The date you requested the quota increase.
- Last updated date The last date the request received an update.

View details about a service, quota name, and status in the **Quota request history** table by choosing one of the entries.

Using AWS CLI

To view the quota request history

 The ListRequestedServiceQuotaChangeHistory operation, which submits the request, requires a QuotaRequestedAtLevel parameter. The following CLI example is for all resource and account level requests.

```
$ aws servicequotas list-requested-service-quota-change-history \
        --quota-applied-at-level ALL
{
    "RequestedQuotas": [
            "QuotaName": "Instances per domain",
        "Status": "PENDING",
        "DesiredValue": 200.0,
        "Created": 1580446904.067,
        "QuotaArn": "arn:aws:servicequotas:us-east-1:123456789012:es/
L-6408ABDE",
        "ServiceName": "Amazon OpenSearch Service",
        "GlobalQuota": false,
        "ServiceCode": "es",
        "QuotaCode": "L-6408ABDE",
        "Requester": "{\"accountId\":\"123456789012\",\"callerArn\":
\"arn:aws:iam::123456789012:root\"}",
        "QuotaRequestedAtLevel": "RESOURCE",
        "Id": "a12345",
        "Unit": "None"
        "QuotaContext": {
                         "ContextId": "arn:aws:es:us-east-1:123456789012:domain/
opensearch-domain-1"
                         "ContextScopeType": "AWS::OpenSearchService::Domain",
```

Verifying your quota request 20

```
"ContextScope": "RESOURCE"
        }
      },
      {
        "QuotaName": "Instances per domain",
        "Status": "PENDING",
        "DesiredValue": 200.0,
        "Created": 1580446904.067,
        "QuotaArn": "arn:aws:servicequotas:us-east-1:123456789012:es/
L-6408ABDE",
        "ServiceName": "Amazon OpenSearch Service",
        "GlobalQuota": false,
        "ServiceCode": "es",
        "QuotaCode": "L-6408ABDE",
        "Requester": "{\"accountId\":\"123456789012\",\"callerArn\":
\"arn:aws:iam::123456789012:root\"}",
        "QuotaRequestedAtLevel": "RESOURCE",
        "Id": "a12345",
        "Unit": "None"
        "QuotaContext": {
                         "ContextId": "arn:aws:es:us-east-1:123456789012:domain/
opensearch-domain-2",
                         "ContextScopeType": "AWS::OpenSearchService::Domain",
                         "ContextScope": "RESOURCE"
        }
      },
      {
            "QuotaName": "Domains per Region",
            "Status": "PENDING",
            "DesiredValue": 120.0,
            "Created": 1580446904.067,
            "Adjustable": true,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-076D529E",
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-076D529E",
            "Requester": "{\"accountId\":\"123456789012\",\"callerArn\":
\"arn:aws:iam::123456789012:root\"}"
            "QuotaRequestedAtLevel": "ACCOUNT",
            "Id": "a123456",
            "Unit": "None",
      },
```

Verifying your quota request 21

```
{
            "QuotaName": "Instances per domain",
            "Status": "PENDING"
            "DesiredValue": 300.0,
            "Created": 1580446904.067,
            "QuotaArn": "arn:aws:servicequotas:us-east-1::123456789012:es/
L-6408ABDE",
            "ServiceName": "Amazon OpenSearch Service",
            "GlobalQuota": false,
            "ServiceCode": "es",
            "QuotaCode": "L-6408ABDE",
            "Requester": "{\"accountId\":\"123456789012\",\"callerArn\":
\"arn:aws:iam::123456789012:root\"}"
            "QuotaRequestedAtLevel": "ACCOUNT",
            "Id": "a1234567",
            "Unit": "None",
            "QuotaContext": {
                             "ContextId": "*",
                             "ContextScopeType":
 "AWS::OpenSearchService::Domain",
                             "ContextScope": "RESOURCE"
                         }
      }
    ]
}
```

Verifying your quota request 22

Managing Service Quotas resources with tags

You can use tags to categorize resources by purpose, owner, environment, or other criteria. A *tag* is a custom attribute label that you add to an AWS resource to make it easier to identify, organize, and search for resources. Each tag includes two parts:

- A tag key, such as CostCenter, Environment, or Project. Tag keys are case sensitive.
- A **tag value**, such as 111122223333 or Production. You can set the value of a tag to an empty string, but you can't set the value of a tag to null. Omitting the tag value is the same as using an empty string. Like tag keys, tag values are case sensitive.

Tags help you do the following:

- Identify and organize your AWS resources. Many services support tagging, so you can assign the same tag to resources from different services to indicate that the resources are related.
- Track your AWS costs. You activate these tags on the AWS Billing and Cost Management
 dashboard. AWS uses the tags to categorize your costs and deliver a monthly cost allocation
 report to you. For more information, see Use cost allocation tags in the AWS Billing User Guide.
- Control access to your AWS resources. For more information, see <u>Controlling access using tags</u> in the *IAM User Guide*.

Topics

- Service Quotas resources that support tagging
- Tag restrictions
- Enabling the required permissions for tagging Service Quotas resources
- Managing Service Quotas tags
- Controlling access using Service Quotas tags

Service Quotas resources that support tagging

Service Quotas supports tagging the **Applied quotas** resource. An applied quota is a quota that you requested an increase for, *and* the increase was approved by AWS Support.

Supported resources 23

Important

You can only tag quotas if they have an applied quota value. Quotas with default quota values cannot be tagged.

Do not store personally identifiable information (PII) or other confidential or sensitive information in tags. Tags are not intended to be used for private or sensitive data.

Tag restrictions

Restrictions apply to tags on Service Quotas resources, including:

- Maximum number of tags that you can assign to a resource 50
- Maximum tag key length 128 Unicode characters
- Maximum tag value length 256 Unicode characters
- Valid characters for tag key and value a-z, A-Z, 0-9, space, and the following characters: _ . : / = + - and @
- Tag keys and values are case sensitive.
- Don't use aws: as a prefix for tag keys. It is reserved for AWS use.

Enabling the required permissions for tagging Service Quotas resources

You must configure permissions to allow your users or roles to manage tags in Service Quotas. The permissions that are required to administer tags generally correspond to the API operations for the task.

To allow IAM principles, such as roles or users, to use Service Quotas for tagging operations, attach the ServiceQuotasReadOnlyAccessAWS managed policy to the principals.

Task	Required permission
Add tags to applied quotas	servicequotas:List
	TagsForResource

Tag restrictions

Task	Required permission
	servicequotas:TagR esource
View tags for an applied quota	servicequotas:List TagsForResource
Remove existing tags from an applied quota	servicequotas:Unta gResource
Edit existing tag values for applied quotas	servicequotas:List TagsForResource
	servicequotas:TagR esource
	servicequotas:Unta gResource

Managing Service Quotas tags

You can manage Service Quotas tags by using the AWS Management Console, the AWS CLI, or the AWS API.

Managing tags from the AWS Management Console

- 1. Sign in to the AWS Management Console and open the Service Quotas console at https://console.aws.amazon.com/servicequotas/home.
- 2. In the navigation page, choose AWS services.
- 3. Choose an AWS service from the list, or enter the name of the service in the search box.
- 4. Choose a service that has a value in the **Applied quota value** column.
- 5. In the **Tags** section, choose **Manage tags**. This option is not available for quotas that don't have an applied quota value.

Managing tags 25

6. You can add or remove tags, or you can edit tag values for existing tags. Enter a name for the tag in **Key**. You can add an optional value for the tag in **Value**.

7. After making all of your changes to tags, choose **Save changes**.

If the operation is successful, you return to the quota details page where you can verify your changes. If the operation fails, follow the instructions in the error message to resolve it.

Managing Service Quotas tags using the AWS CLI or API

To manage Service Quotas tags using the CLI or API, choose a management task and use the corresponding CLI command or API call.

Tag management task	CLI command	API call
Add tags to applied quotas	aws service-quotas <pre>tag-resource</pre>	TagResource
View tags for an applied quota	aws service-quotas <u>list-tags-for-reso</u> <u>urce</u>	ListTagsForResource
Delete existing tag values for applied quotas	aws service-quotas untag-resource	<u>UntagResource</u>

Controlling access using Service Quotas tags

To control access to Service Quotas resources based on tags, you provide the tag information in the <u>condition element</u> of a policy using the aws:ResourceTag/key-name, aws:RequestTag/key-name, or aws:TagKeys condition keys. For more information about these condition keys, see <u>Controlling access to AWS resources using resource tags</u> in the *IAM User Guide*.

For example, when you attach the following policy to an AWS Identity and Access Management (IAM) role or user, that principal can request an increase to Amazon Athena applied quotas that are tagged with the tag key **Owner** and tag value **admin**.

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

Managing tags (CLI or API) 26

You can also attach tags to IAM principals to use attribute-based access control (ABAC). ABAC is an authorization strategy that defines permissions based on attributes. Tagging entities and resources is the first step of ABAC. Then you design ABAC policies to allow operations when the principal's tag matches the tag on the resource that they're trying to access. ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

For more information about ABAC, see <u>What is ABAC?</u> in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see <u>IAM tutorial</u>: <u>Define permissions to access AWS resources based on tags</u> in the *IAM User Guide*.

Controlling access using tags 27

Using Service Quotas request templates



Note

You can use quota request templates only with AWS accounts that are members of an organization managed by AWS Organizations.

A quota request template helps you save time when customizing quotas for new AWS accounts in your organization. To use a template, configure the desired service quota increases for new accounts. Then, enable template association. This associates the template with your organization in AWS Organizations. Whenever new accounts are created in your organization, the template automatically requests quota increases for you.

To use a request template, you must use AWS Organizations and the new accounts must be created in the same organization. Your organization must have all features enabled, all features. If you use consolidated billing features only, you can't use quota request templates.

You can update the request template by adding or removing service quotas. You can also increase the values for adjustable quotas. As soon as you adjust the template, those service quota values are requested for new accounts. Updating a request template doesn't update quota values for existing accounts.

To enable template association

- Sign in to the AWS Management Console and open the Service Quotas console at https:// console.aws.amazon.com/servicequotas/home.
- 2. In the navigation pane, expand **Organization**, and then choose **Quota request template**.
- 3. In the **Template association** section, choose **Enable**.

To add a quota to your request template

- Sign in to the AWS Management Console and open the Service Quotas console at https:// console.aws.amazon.com/servicequotas/home.
- 2. In the navigation pane, expand **Organization**, and then choose **Quota request template**.
- In the **Added quotas** section, choose **Add quota**.



Note

You add up to 10 quotas to your request template.

On the Add quota page, choose a Region, Service, Quota, and Desired quota value, and then choose Add.

To remove a quota from your request template

You can remove service quota requests from the template regardless of whether the template is associated with an organization. If you reach the maximum number of service quota requests, you might need to remove some quotas from your request template.

- Sign in to the AWS Management Console and open the Service Quotas console at https:// 1. console.aws.amazon.com/servicequotas/home.
- In the navigation pane, expand **Organization**, and then choose **Quota request template**. 2.
- 3. In the **Added quotas** section, select the option button for the quota that you want to remove.
- Choose Remove. 4.

To disable the template association

If you disable the automatic template association, new accounts receive the AWS default quota values for all quotas. Disabling the template association from the organization doesn't delete the service quota requests from the template. You can continue to edit the service quotas in the template.

- Sign in to the AWS Management Console and open the Service Quotas console at https:// 1. console.aws.amazon.com/servicequotas/home.
- 2. In the navigation pane, expand **Organization**, and then choose **Quota request template**.
- 3. In the **Template association** section, choose **Disable**.

Security in Service Quotas

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The <u>shared responsibility model</u> describes this as security *of* the cloud and security *in* the cloud:

- Security of the cloud AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the <u>AWS</u>
 <u>Compliance Programs</u>. To learn about the compliance programs that apply to Service Quotas, see AWS Services in Scope by Compliance Program.
- **Security in the cloud** Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Service Quotas. The following topics show you how to configure Service Quotas to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Service Quotas resources.

Contents

- Data protection in Service Quotas
- Logging and monitoring Service Quotas
- Identity and access management in Service Quotas
- Compliance validation for Service Quotas
- Resilience in Service Quotas
- Infrastructure Security in Service Quotas

Data protection in Service Quotas

The AWS <u>shared responsibility model</u> applies to data protection in Service Quotas. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the

Data protection 30

AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the Data Privacy FAQ. For information about data protection in Europe, see the AWS Security Blog.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-2 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see Federal Information Processing Standard (FIPS) 140-2.

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Service Quotas or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Encryption at rest

Quota-related information, such as applied quota values and quota request history, are encrypted at rest. Unique encryption keys are used for each Region.

Encryption in transit

Customer requests and all associated data is encrypted in transit using <u>Transport Later Security</u> (TLS) 1.2 or later. All Service Quotas endpoints support HTTPS for encrypting data in transit.

Encryption at rest 31

Logging and monitoring Service Quotas

Overview

Monitoring is an important part of maintaining the reliability, availability, and performance of Service Quotas and your other AWS solutions. AWS provides the following monitoring tools to watch Service Quotas, report when something is wrong, and take automatic actions when appropriate:

- AWS CloudTrail captures API calls and related events made by or on behalf of your AWS account and delivers the log files to an Amazon S3 bucket that you specify. You can identify which users and accounts called AWS, the source IP address from which the calls were made, and when the calls occurred. For more information, see the AWS CloudTrail User Guide.
- Amazon CloudWatch monitors your AWS resources and the applications you run on AWS in real
 time. You can collect and track metrics, create customized dashboards, and set alarms that notify
 you or take actions when a specified metric reaches a threshold that you specify. For example,
 you can have CloudWatch track CPU usage or other metrics of your Amazon EC2 instances
 and automatically launch new instances when needed. For more information, see the Amazon CloudWatch User Guide.

Logging Service Quotas API calls using AWS CloudTrail

Service Quotas is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Service Quotas. CloudTrail captures all API calls for Service Quotas as events. The calls captured include calls from the Service Quotas console and code calls to the Service Quotas API operations. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Service Quotas. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Service Quotas, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, see the AWS CloudTrail User Guide.

Service Quotas information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When activity occurs in Service Quotas, that activity is recorded in a CloudTrail event along with other AWS service events

Logging and monitoring 32

in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see Viewing events with CloudTrail Event history.

For an ongoing record of events in your AWS account, including events for Service Quotas, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- · Overview for creating a trail
- CloudTrail supported services and integrations
- Configuring Amazon SNS notifications for CloudTrail
- Receiving CloudTrail log files from multiple regions and Receiving CloudTrail log files from multiple accounts

All Service Quotas actions are logged by CloudTrail and are documented in the <u>Service Quotas API Reference</u>. For example, calls to the GetServiceQuota, RequestServiceQuotaIncrease and ListAWSDefaultServiceQuotas actions generate entries in the CloudTrail log files.

Every event or log entry contains information that helps you determine who made the request.

- AWS account root credentials.
- Temporary security credentials from an AWS Identity and Access Management role or federated user.
- Long-term security credentials from an IAM user.
- Another AWS service.

For more information, see the <u>CloudTrail userIdentity element</u>.

Understanding Service Quotas log file entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the RequestQuotaIncrease action.

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "IAMUser",
        "principalId": "AIDA123456789012Example",
        "arn": "arn:aws:iam::123456789012:user/admin",
        "accountId": "123456789012",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "userName": " admin",
        "sessionContext": {
            "sessionIssuer": {},
            "webIdFederationData": {},
            "attributes": {
                "creationDate": "2022-01-24T16:57:04Z",
                "mfaAuthenticated": "true"
            }
        }
    },
    "eventTime": "2022-01-24T17:00:15Z",
    "eventSource": "servicequotas.amazonaws.com",
    "eventName": "RequestServiceQuotaIncrease",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "172.21.16.1",
    "userAgent": "aws-internal/3 aws-sdk-java/1.12.127
Linux/5.4.147-83.259.amzn2int.x86_64 OpenJDK_64-Bit_Server_VM/25.312-b07
java/1.8.0_312 vendor/Oracle_Corporation cfg/retry-mode/standard",
    "requestParameters": {
        "serviceCode": "ec2",
        "quotaCode": "L-CEED54BB",
        "desiredValue": 10
    },
    "responseElements": {
        "requestedQuota": {
            "id": "cd3ad3d9-2776-4ef1-a904-4c229d1642ee",
            "serviceCode": "ec2",
            "serviceName": "Amazon Elastic Compute Cloud (Amazon EC2)",
            "quotaCode": "L-CEED54BB",
            "quotaName": "EC2-Classic Elastic IPs",
            "desiredValue": 10,
            "status": "PENDING",
```

```
"created": "Jan 24, 2022 5:00:15 PM",
            "requester": "{\"accountId\":\"123456789012\",\"callerArn\":
\"arn:aws:iam::123456789012:user/admin\"}",
            "quotaArn": "arn:aws:servicequotas:us-east-1:123456789012:ec2/L-CEED54BB",
            "globalQuota": false,
            "unit": "None"
        }
    },
    "requestID": "3d3f5cdc-af30-4121-b69a-84b2f5c33be5",
    "eventID": "0cb51588-e460-4e00-bc48-a9d4820cad83",
    "readOnly": false,
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "123456789012",
    "eventCategory": "Management"
}
```

This example shows that the user named admin generated a request for additional Amazon Elastic Compute Cloud Elastic IP addresses on January 24, 2022. The requested increase was 10, an increase of 5 from the default quota of 5.

The following is an example of an approved quota increase in Service Quotas:

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "accountId": "123456789012",
        "invokedBy": "servicequotas.amazonaws.com"
    },
    "eventTime": "2022-01-24T17:02:17Z",
    "eventSource": "servicequotas.amazonaws.com",
    "eventName": "UpdateServiceQuotaIncreaseRequestStatus",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "servicequotas.amazonaws.com",
    "userAgent": "servicequotas.amazonaws.com",
    "requestParameters": null,
    "responseElements": null,
    "eventID": "e331b0a0-9395-4895-aeba-73cbab9ebcb0",
    "readOnly": false,
    "eventType": "AwsServiceEvent",
    "managementEvent": true,
    "recipientAccountId": "123456789012",
    "serviceEventDetails": {
```

```
"requestId": "cdc5f1f78739459e6642407bb2bZK08GKUM",
    "newStatus": "CASE_CLOSED",
    "createTime": "2022-01-24T17:00:15.363Z",
    "newQuotaValue": "10.0",
    "serviceName": "Amazon Elastic Compute Cloud (Amazon EC2)",
    "quotaName": "EC2-Classic Elastic IPs",
    "unit": "None"
    },
    "eventCategory": "Management"
}
```

From the serviceEventDetails section, you can determine that AWS Support approved the request for a quota increase to 10 Elastic IP addresses, and closed the request. The newQuotaValue displays 10 as the new quota.

Service Quotas and Amazon CloudWatch alarms

You can create Amazon CloudWatch alarms to notify you when you're close to a quota value threshold. Setting an alarm can help alert you if you need to request a quota increase.

To create a CloudWatch alarm for a quota

- 1. Sign in to the AWS Management Console and open the Service Quotas console at https://console.aws.amazon.com/servicequotas/home.
- 2. In the navigation pane, choose **AWS services** and then select a service.
- 3. Select a guota that supports CloudWatch alarms.

If you actively use the quota, utilization appears beneath the quota description. If CloudWatch alarms are supported, the CloudWatch alarms section appears at the bottom of the page.

- 4. In Amazon CloudWatch alarms, choose Create.
- 5. For **Alarm threshold**, choose a threshold.
- For Alarm name, enter a name for the alarm. This name must be unique within the AWS account.
- 7. Choose Create.

Using CloudWatch alarms 36



Note

To add a notification to the CloudWatch alarm, see Creating a CloudWatch alarm based on a static threshold in the Amazon CloudWatch User Guide.

To delete a CloudWatch alarm

- 1. Choose the service quota with the alarm.
- 2. Select the alarm.
- Choose Delete. 3.

Identity and access management in Service Quotas

AWS uses security credentials to identify you and to grant you access to your AWS resources. You can use features of AWS Identity and Access Management (IAM) to allow other users, services, and applications to use your AWS resources fully or in a limited way. You can do this without sharing your security credentials.

By default, principals, such as IAM roles or users, don't have permission to create, view, or modify AWS resources. To allow a principal to access resources such as a load balancer, and to perform tasks, perform the following steps:

- 1. Create an IAM policy that grants the principal permission to use the specific resources and API actions they need.
- 2. Attach the policy to the IAM principal or the group that the principal belongs to.

When you attach a policy to a user or group of users, it allows or denies the users permission to perform the specified tasks on the specified resources.

For example, you can use IAM to create roles or users as the principals in your AWS account. A principal can represent a person, a system, or an application. Then you grant permissions to the principals to perform specific actions on the specified resources using an IAM policy.

Grant permissions using IAM policies

When you attach a policy to a principal or a group of principals, it allows or denies those principals permission to perform the specified tasks on the specified resources.

An IAM policy is a JSON document that consists of one or more statements. Each statement is structured as shown in the following example.

```
{
  "Version": "2012-10-17",
  "Statement":[{
      "Effect": "effect",
      "Action": "action",
      "Resource": "resource-arn",
      "Condition": {
            "condition": {
                "key":"value"
            }
        }
    }
}]
```

- **Effect** The value for **effect** can be either Allow or Deny. By default, IAM principals don't have permission to use resources and API actions, so all requests are denied. An explicit allow overrides the default. An explicit deny overrides any allows.
- **Action** The value for **action** is the specific API action for which you are granting or denying permission. For more information about specifying Action, see API actions for Service Quotas.
- Resource The resource that's affected by the action. With some Service Quotas API actions, you can restrict the permissions granted or denied to a specific quota. To do so, specify its Amazon Resource Name (ARN) in this statement. Otherwise, you can use the wildcard character (*) to specify all Service Quotas resources. For more information, see Service Quotas resources.
- Condition You can optionally use conditions to control when your policy is in effect. For more information, see Condition keys for Service Quotas.

For more information, see the <u>IAM User Guide</u>.

AWS managed policies for Service Quotas

An AWS managed policy is a standalone policy that is created and administered by AWS. AWS managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that AWS managed policies might not grant least-privilege permissions for your specific use cases because they're available for all AWS customers to use. We recommend that you reduce permissions further by defining customer managed policies that are specific to your use cases.

You cannot change the permissions defined in AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. AWS is most likely to update an AWS managed policy when a new AWS service is launched or new API operations become available for existing services.

For more information, see AWS managed policies in the IAM User Guide.

AWS managed policy: ServiceQuotasFullAccess

You can attach ServiceQuotasFullAccess to your users, groups, and roles.

This policy grants permissions that allow full administrative control of the Service Quotas service. You can perform all tasks involved in viewing and managing your quotas for AWS services in Service Quotas in the AWS Regions in your account.

Permissions details

This policy includes permissions that allow all actions for Service Quotas, including viewing AWS default values and applied values, requesting a service quota increase, and viewing current utilization of resources. This policy also includes 18 permissions that are not part of Service Quotas and can be broadly split into **non-mutating** and **mutating** operations. Non-mutating operations include permissions from trusted advisors to retrieve applied quota value and view current utilization of resources. Mutating operations include permission to create and delete alarms on utilization of resources, and permissions to create the service-linked role necessary to create a support case on your behalf while requesting a quota increase.

This policy includes the following non-mutating and mutating operations that are *not* part of Service Quotas:

Non-mutating operations

• autoscaling: DescribeAccountLimits – Allows Service Quotas to retrieve applied quota value for AWS Auto Scaling quotas.

• cloudformation: DescribeAccountLimits – Allows Service Quotas to retrieve applied quota value for AWS CloudFormation quotas.

- cloudwatch: DescribeAlarmsForMetric Allows you to retrieve alarms for specified metrics from Service Quotas that were created for notifying automatically whenever a specified quota reaches a percentage of the maximum or reaches the maximum level.
- cloudwatch: DescribeAlarms Allows you to retrieve alarms from Service Quotas that were created for notifying automatically whenever a specified quota reaches a percentage of the maximum or reaches the maximum level.
- cloudwatch: GetMetricData Allows Service Quotas to view current utilization of resources.
- cloudwatch: GetMetricStatistics Allows Service Quotas to view current utilization of resources.
- dynamodb: DescribeLimits Allows Service Quotas to retrieve applied quota value for DynamoDB quotas.
- elasticloadbalancing: DescribeAccountLimits Allows Service Quotas to retrieve applied quota value for Elastic Load Balancing quotas.
- iam: GetAccountSummary Allows Service Quotas to retrieve applied quota value for IAM.
- kinesis:DescribeLimits Allows Service Quotas to retrieve applied quota value for Amazon Kinesis quotas.
- organizations:DescribeAccount and organizations:DescribeOrganization –
 Allows Service Quotas to create and execute quota templates.
- rds:DesceibeAccountAttributes Allows Service Quotas to retrieve applied quota value for Amazon RDS quotas.
- route53:GetAccountLimit Allows Service Quotas to retrieve applied quota value for Amazon Route 53 quotas.
- tag:GetTagKeys Allows Service Quotas to get tag keys currently in use in the specified AWS Region for the calling account.
- tag: GetTagValues Allows Service Quotas to get tag values for the specified key that are used in the specified AWS Region for the calling account.

Mutating operations

• cloudwatch: PutMetricAlarm – Allows Service Quotas to create an alarm for notifying you automatically whenever a specified quota reaches a percentage of the maximum or the maximum level.

- cloudwatch:DeleteAlarms Allows Service Quotas to delete the specified alarm.
- organizations: EnableAWSServiceAccess Allows Service Quotas to create a servicelinked role in all the accounts in your organization. This allows Service Quotas to perform operations on your behalf in your organization and its accounts.

 iam:CreateServicelinkedRole – Allows Service Quotas to create an IAM role that allows Service Quotas to create a support case on your behalf when you request a quota increase.

To see the latest version of this AWS managed policy, see ServiceQuotasFullAccess in the AWS Managed Policy Reference Guide.

AWS managed policy: ServiceQuotasReadOnlyAccess

You can attach ServiceQuotasReadOnlyAccess to your users, groups, and roles.

This policy grants permissions that allow users to view their AWS default quotas, applied quotas, and view current utilization of resources.

Permissions details

This policy includes permissions that allow your to perform the Service Quotas Get*, and List* operations to view your AWS default quotas and applied quotas. You can also view current utilization of resources.



Note

This policy does not allow you to request a service quota increase.

This policy includes the following non-mutating operations that are *not* part of Service Quotas:

Non-mutating operations

- autoscaling: DescribeAccountLimits Allows Service Quotas to retrieve applied quota value for AWS Auto Scaling quotas.
- cloudformation: DescribeAccountLimits Allows Service Quotas to retrieve applied quota value for AWS CloudFormation quotas.

• cloudwatch: DescribeAlarmsForMetric – Allows you to retrieve alarms for specified metrics from Service Quotas that were created for notifying automatically whenever a specified quota reaches a percentage of the maximum or reaches the maximum level.

- cloudwatch: DescribeAlarms Allows you to retrieve alarms from Service Quotas that were created for notifying automatically whenever a specified quota reaches a percentage of the maximum or reaches the maximum level.
- cloudwatch: GetMetricData Allows Service Quotas to view current utilization of resources.
- cloudwatch: GetMetricStatistics Allows Service Quotas to view current utilization of resources.
- dynamodb: DescribeLimits Allows Service Quotas to retrieve applied quota value for DynamoDB quotas.
- elasticloadbalancing: DescribeAccountLimits Allows Service Quotas to retrieve applied quota value for Elastic Load Balancing quotas.
- iam: GetAccountSummary Allows Service Quotas to retrieve applied quota value for IAM.
- kinesis:DescribeLimits Allows Service Quotas to retrieve applied quota value for Amazon Kinesis quotas.
- organizations:DescribeAccount and organizations:DescribeOrganization –
 Allows Service Quotas to create and execute quota templates.
- rds:DesceibeAccountAttributes Allows Service Quotas to retrieve applied quota value for Amazon RDS quotas.
- route53:GetAccountLimit Allows Service Quotas to retrieve applied quota value for Amazon Route 53 quotas.
- tag:GetTagKeys Allows Service Quotas to get tag keys currently in use in the specified AWS Region for the calling account.
- tag: GetTagValues Allows Service Quotas to get tag values for the specified key that are used in the specified AWS Region for the calling account.

To see the latest version of this AWS managed policy, see ServiceQuotasReadOnlyAccess in the AWS Managed Policy Reference Guide.

AWS managed policy: ServiceQuotasServiceRolePolicy

This policy is attached to a service-linked role that allows the service to perform actions on your behalf. You cannot attach this policy to your users, groups, or roles.

This policy grants permissions that allows Service Quotas to create support cases on your behalf.

Permissions details

This policy includes the following operations:

- support: CreateCase Allows Service Quotas to create support cases on your behalf when you request a quota increase.
- support:DescribeCases Allows Service Quotas to retrieve the details and status of your support case for the quota increase request.
- support: RresolveCase Allows Service Quotas to resolve support cases on your behalf.

To see the latest version of this AWS managed policy, see ServiceQuotasServiceRolePolicy in the AWS Managed Policy Reference Guide.

Service Quotas updates to AWS managed policies

View details about updates to AWS managed policies for Service Quotas since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the Service Quotas Document history page.

Change	Description	Date
ServiceQuotasFullAccess – New policy	Added a new AWS managed policy that allows full administrative control of the Service Quotas service. You can perform all tasks involved in viewing and managing your quotas for AWS services in Service Quotas in the AWS Regions in your account.	May 30, 2024

Change	Description	Date
ServiceQuotasReadO nlyAccess – New policy	Added a new AWS managed policy that allows users to view their AWS default quotas, applied quotas, and view current utilization of resources.	May 30, 2024
ServiceQuotasServiceRolePol icy – New policy	Added a new AWS managed policy that allows Service Quotas to create support cases on your behalf.	May 30, 2024
Service Quotas started tracking changes	Service Quotas started tracking changes for its AWS managed policies.	May 30, 2024

API actions for Service Quotas

In the Action element of your IAM policy statement, you can specify any API action that Service Quotas offers. You must prefix the action name with the lowercase string servicequotas:, as shown in the following example.

```
"Action": "servicequotas:GetServiceQuota"
```

To specify multiple actions in a single statement, enclose them in square brackets and separate them with a comma, as shown in the following example.

```
"Action": [
    "servicequotas:ListRequestedServiceQuotaChangeHistory",
    "servicequotas:ListRequestedServiceQuotaChangeHistoryByQuota"
]
```

You can also specify multiple actions using the wildcard character (*). The following example specifies all API action names for Service Quotas that start with Get.

```
"Action": "servicequotas:Get*"
```

To specify all API actions for Service Quotas, use the wildcard character (*), as shown in the following example.

```
"Action": "servicequotas:*"
```

For the list of API actions for Service Quotas, see Service Quotas Actions.

Service Quotas resources

Resource-level permissions refers to the ability to specify which resources users are allowed to perform actions on. For API actions that support resource-level permissions, you can control the resources that users are allowed to use with the action. To specify a resource in a policy statement, you must use its Amazon Resource Name (ARN).

The ARN for a quota has the format shown in the following example.

```
arn:aws:servicequotas:region-code:account-id:service-code/quota-code
```

For API actions that don't support resource-level permissions, you must specify the resource statement shown in the following example.

```
"Resource": "*"
```

Resource-level permissions for Service Quotas

The following Service Quotas actions support resource-level permissions:

- PutServiceQuotaIncreaseRequestIntoTemplate
- RequestServiceQuotaIncrease

For more information, see Actions defined by Service Quotas in the Service Authorization Reference.

Condition keys for Service Quotas

When you create a policy, you can specify the conditions that control when the policy is in effect. Each condition contains one or more key-value pairs. There are global condition keys and service-specific condition keys.

Service Quotas resources 45

The service quotas: service key is specific to Service Quotas. The following Service Quotas API actions support this key:

- PutServiceQuotaIncreaseRequestIntoTemplate
- RequestServiceQuotaIncrease

For more information about global condition keys, see <u>AWS Global Condition Context Keys</u> in the *IAM User Guide*.

Predefined AWS managed policies for Service Quotas

The managed policies created by AWS grant the required permissions for common use cases. You can attach these policies to your IAM principals, based on the access to Service Quotas that they require:

- ServiceQuotasFullAccess Grants full access required to use Service Quotas features.
- ServiceQuotasReadOnlyAccess Grants read-only access to Service Quotas features.

Compliance validation for Service Quotas

To learn whether an AWS service is within the scope of specific compliance programs, see <u>AWS</u> <u>services in Scope by Compliance Program</u> and choose the compliance program that you are interested in. For general information, see <u>AWS Compliance Programs</u>.

You can download third-party audit reports using AWS Artifact. For more information, see Downloading Reports in AWS Artifact.

Your compliance responsibility when using AWS services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- <u>Security and Compliance Quick Start Guides</u> These deployment guides discuss architectural
 considerations and provide steps for deploying baseline environments on AWS that are security
 and compliance focused.
- Architecting for HIPAA Security and Compliance on Amazon Web Services This whitepaper describes how companies can use AWS to create HIPAA-eligible applications.



Note

Not all AWS services are HIPAA eligible. For more information, see the HIPAA Eligible Services Reference.

- AWS Compliance Resources This collection of workbooks and guides might apply to your industry and location.
- AWS Customer Compliance Guides Understand the shared responsibility model through the lens of compliance. The guides summarize the best practices for securing AWS services and map the guidance to security controls across multiple frameworks (including National Institute of Standards and Technology (NIST), Payment Card Industry Security Standards Council (PCI), and International Organization for Standardization (ISO)).
- Evaluating Resources with Rules in the AWS Config Developer Guide The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- AWS Security Hub This AWS service provides a comprehensive view of your security state within AWS. Security Hub uses security controls to evaluate your AWS resources and to check your compliance against security industry standards and best practices. For a list of supported services and controls, see Security Hub controls reference.
- Amazon GuardDuty This AWS service detects potential threats to your AWS accounts, workloads, containers, and data by monitoring your environment for suspicious and malicious activities. GuardDuty can help you address various compliance requirements, like PCI DSS, by meeting intrusion detection requirements mandated by certain compliance frameworks.
- AWS Audit Manager This AWS service helps you continuously audit your AWS usage to simplify how you manage risk and compliance with regulations and industry standards.

Resilience in Service Quotas

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

Resilience 47

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure.

In addition to the AWS global infrastructure, Service Quotas offers several features to help support your data resiliency and backup needs.

Infrastructure Security in Service Quotas

As a managed service, Service Quotas is protected by AWS global network security. For information about AWS security services and how AWS protects infrastructure, see <u>AWS Cloud Security</u>. To design your AWS environment using the best practices for infrastructure security, see <u>Infrastructure</u> <u>Protection</u> in *Security Pillar AWS Well-Architected Framework*.

You use AWS published API calls to access Service Quotas through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the <u>AWS Security Token Service</u> (AWS STS) to generate temporary security credentials to sign requests.

You can call these API operations from any network location, but Service Quotas does support resource-based access policies, which can include restrictions based on the source IP address.

Infrastructure Security 48

Service quotas for Service Quotas

The following tables list the maximum values for Service Quotas resources for your AWS account.

All of these quota values are per AWS Region, unless noted otherwise.

You can't adjust these quota values.

Increase requests

Quota	Default
Active service quota increase requests per account	20
Active service quota increase requests per Region	2
Active service quota increase requests per quota	1
Active service quota increase requests per resource	1

API request rates

Quota	Default
GetAWSDefaultServiceQuota requests per second	5
Additional GetAWSDefaultServiceQuota requests per second sent in one burst	5
GetRequestedServiceQuotaChange requests per second	5
Additional GetRequestedServiceQuotaChange requests per second sent in one burst	5
GetServiceQuota requests per second	5
Additional GetServiceQuota requests per second sent in one burst	5
ListAWSDefaultServiceQuotas requests per second	10

Quota	Default
Additional ListAWSDefaultServiceQuotas requests per second sent in one burst	10
ListRequestedServiceQuotaChangeHistory requests per second	5
Additional ListRequestedServiceQuotaChangeHistory requests per second sent in one burst	5
ListRequestedServiceQuotaChangeHistoryByQuota requests per second	5
Additional ListRequestedServiceQuotaChangeHisto ryByQuota requests per second sent in one burst	5
ListServiceQuotas requests per second	10
Additional ListServiceQuotas requests per second sent in one burst	10
ListServices requests per second	10
Additional ListServices requests per second sent in one burst	10
ListTagsForResource requests per second	10
ListTagsForResource requests per second sent in one burst	10
RequestServiceQuotaIncrease requests per second	3
Additional RequestServiceQuotaIncrease requests per second sent in one burst	3
TagResource requests per second	10
TagResource requests per second sent in one burst	10
UntagResource requests per second	10

Quota	Default
UntagResource requests per second sent in one burst	10

Quota request template API request rates

Quota	Default
AssociateQuotaTemplate requests per second	1
Additional AssociateQuotaTemplate requests per second sent in one burst	1
DeleteServiceQuotaIncreaseRequestFromTemplate requests per second	2
Additional DeleteServiceQuotaIncreaseRequestFro mTemplate requests per second sent in one burst	1
DisassociateQuotaTemplate requests per second	1
Additional DisassociateQuotaTemplate requests per second sent in one burst	1
GetAssociationForQuotaTemplate requests per second	2
Additional GetAssociationForQuotaTemplate requests per second sent in one burst	2
GetServiceQuotaIncreaseRequestFromTemplate requests per second	2
Additional GetServiceQuotaIncreaseRequestFromTe mplate requests per second sent in one burst	1
ListServiceQuotaIncreaseRequestsInTemplate requests per second	2

Quota	Default
Additional ListServiceQuotaIncreaseRequestsInTe mplate requests per second sent in one burst	1
PutServiceQuotaIncreaseRequestIntoTemplate requests per second	1
Additional PutServiceQuotaIncreaseRequestIntoTe mplate per second sent in one burst	1

Overall API request rate

Quota	Default
API requests per AWS account (requests per second)	50
API requests per AWS Organization (requests per second)	100

Service Quotas Document history

The following table describes the important changes to the documentation since the last release of Service Quotas. For notification about updates to this documentation, you can subscribe to an RSS feed.

Change	Description	Date
<u>Updated content</u>	Updated topic titles and reorganized content to improve readability and discoverability.	June 20, 2024
Adding new AWS managed policies	You can now attach ServiceQuotasFullA ccess and ServiceQu otasReadOnlyAccess policies to your users, groups, and roles.	May 30, 2024
Adding support for context based quota management	You now have greater visibilit y and control over your service quotas. View applied values, monitor usage, and programmatically request increases for quotas that not only apply at the AWS account level, but at the resource level.	August 30, 2023
IAM best practices update	Updated guide to align with the IAM best practices . For more information, see Security best practices in IAM.	January 3, 2023
Tagging Service Quotas resources	You can now attach tags to applied quotas and write	December 21, 2020

policies to control access to

those quotas.

Initial release This release introduces June 24, 2019

Service Quotas.