Implementation Guide

Quota Monitor for AWS

Quota Monitor for AWS: Implementation Guide

Table of Contents

Monitor resource usage and send notifications when approaching quotas

Publication date: *September 2016*. Visit the CHANGELOG.md in our GitHub repository to track version-specific improvements and fixes.

The Quota Monitor for AWS solution proactively monitors resource utilization to avoid unexpectedly reaching quota limits. It sends notifications when your Amazon Web Services (AWS) service quotas (previously known as limits) are approaching their maximum value. This solution uses AWS CloudFormation templates to automate the deployment by provisioning the infrastructure resources (also known as the *stack*) automatically.

The solution leverages AWS Trusted Advisor and Service Quotas to monitor resource utilization against quotas for specific AWS services. The solution can send you notifications via email or your existing Slack channel, requesting to increase quotas or to shut down resources before the quota is reached. For more information, refer to Quotas later in this document.

This implementation guide provides an overview of the Quota Monitor for AWS solution, its reference architecture and components, considerations for planning the deployment, configuration steps for deploying the solution to the AWS Cloud. It is intended for solution architects, DevOps engineers, AWS account administrators, and cloud professionals who want to implement Quota Monitor for AWS in their environment.

You can use this navigation table to quickly find answers to these questions:

If you want to . . .	Read . . .
Know the cost for running this solution.	Cost
Understand the security considerations for this solution.	Security
Know how to plan for quotas for this solution.	Quotas
Know which AWS Regions this solution supports.	Supported AWS Regions

If you want to . . .	Read . . .
View or download the AWS CloudFormation templates included in this solution to automatically deploy the infrastructure resources (the "stack") for this solution.	AWS CloudFormation templates
Access the source code and optionally use the AWS Cloud Development Kit (AWS CDK) to deploy the solution.	GitHub repository

Features and benefits

The Quota Monitor for AWS solution provides the following features:

Monitor resource utilization for specific AWS services

The solution leverages AWS Trusted Advisor and Service Quotas to help you monitor resource utilization against quotas for specific AWS services.

Automate Amazon SNS and Slack notifications

The solution publishes alerts to an Amazon Simple Notification Service (Amazon SNS) topic, which you can subscribe to through a notification mechanism of your choice. The solution includes template parameters to configure Amazon SNS notifications to email or an existing Slack channel. Once you receive a notification, you can take corrective measures such as requesting quota increases or shutting down resources.

Choose your deployment scenarios

This solution supports deployment scenarios for both when you are using AWS Organizations and when you are not. For more details, refer to Deployment scenarios.

Start monitoring accounts as they join your organization

When deployed in Organizations mode, the solution uses CloudFormation StackSets to manage template deployments. We configured the StackSets to deploy to accounts that are added to a target organization or organizational units (OUs) within AWS Organizations. This way, you can monitor the new accounts without manual intervention.

Integrate with AWS Service Catalog AppRegistry and Application Manager, a capability of AWS Systems Manager

This solution includes an AWS Service Catalog AppRegistry resource to register the solution's CloudFormation template and its underlying resources as an application in both AWS Service Catalog AppRegistry and Application Manager. With this integration, you can centrally manage the solution's resources and enable application search, reporting, and management actions.

Use cases

Monitoring quotas across your organization or OUs

You can monitor resource utilization across your organization or across different OUs under your AWS Organizations with the Organizations deployment mode in the `quota-monitor-hub.template`.

Monitoring quotas across and outside your organization

You can monitor resource utilization across your organization or across accounts outside your AWS Organizations with the Hybrid deployment mode in the `quota-monitor-hub.template`.

Monitoring quotas within an individual AWS account

You can monitor resource utilization within a single AWS account using the `quota-monitor-hub-no-ou.template`.

For more details, refer to Choose your deployment scenario.

Concepts and definitions

This section describes key concepts and defines terminology specific to this solution.

hub template

AWS CloudFormation template for AWS and all associated components in the primary AWS account that will be used to monitor quotas. See **monitoring account**.

limit

The maximum allocated value for a quota. The term limit was used previously as a synonym for quota.

monitored account

Secondary AWS accounts where a spoke CloudFormation template has been launched to support quota monitoring.

monitoring account

The primary AWS account where the hub CloudFormation template is deployed to monitor secondary AWS accounts.

organization

An entity that you create to consolidate and manage your AWS accounts. An organization has one management account along with zero or more member accounts.

organizational unit

A container for accounts within a root of an organization. An organizational unit (OU) can contain other OUs.

prerequisite template

AWS CloudFormation template used to fulfill the prerequisites needed for monitoring quotas across AWS Organizations.

quota

The maximum number of service resources, actions, operations, and items for your AWS account. Previously known as limit. See quotas.

Service Quotas

An AWS service that helps you manage quotas for many AWS services from one location. Along with looking up the quota values, you can also request a quota increase from the Service Quotas console. See AWS service quotas.

spoke template

AWS CloudFormation template to launch the Quota Monitor for AWS solution and all associated components in secondary accounts to support quota monitoring (monitored accounts).

StackSets

AWS CloudFormation StackSets extend the capability of stacks by enabling you to create, update, or delete stacks across multiple accounts and AWS Regions with a single operation.

For a general reference of AWS terms, refer to the AWS Glossary.

Architecture overview

This section provides a reference implementation architecture diagram for the components deployed with this solution.

Architecture diagram

Deploying this solution with the default parameters deploys the following components in your AWS account.

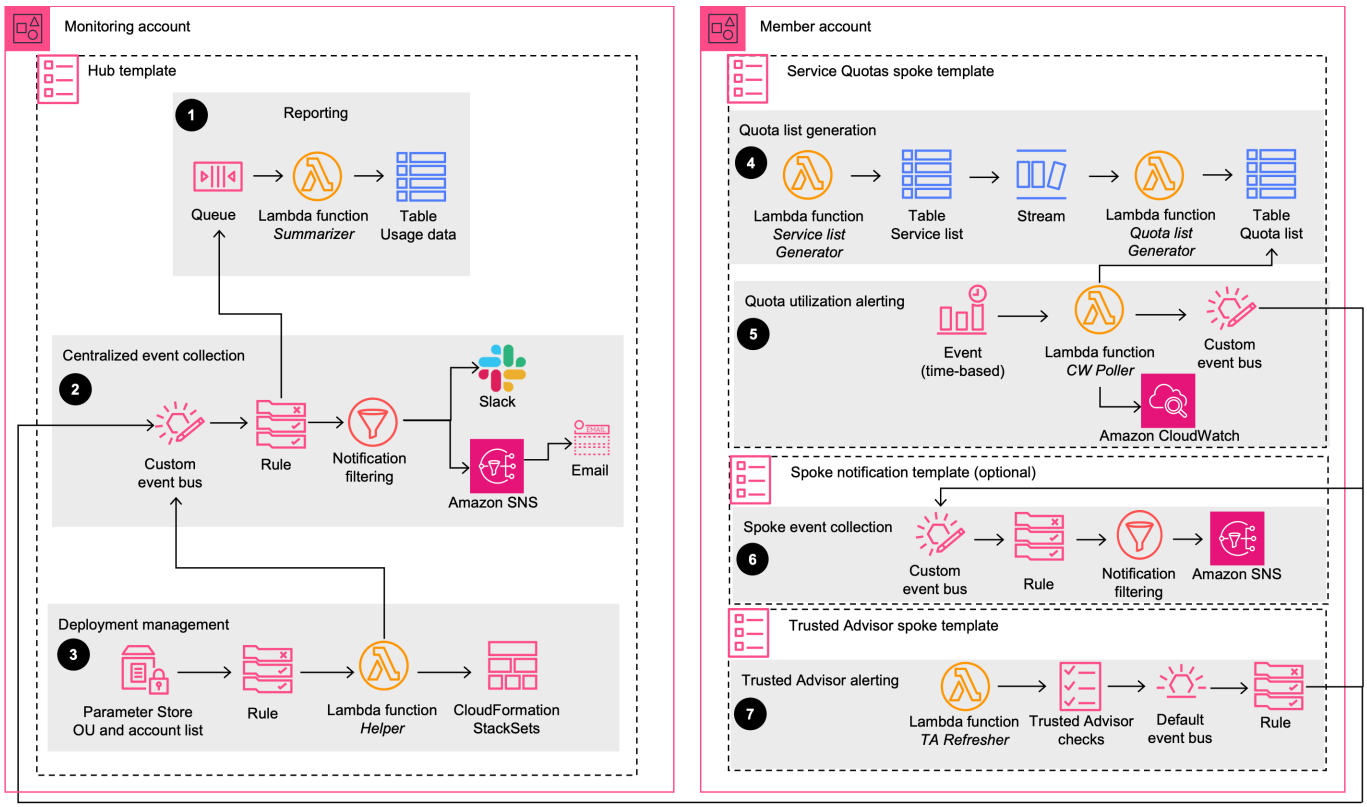

Quota Monitor for AWS includes a **hub template** that you deploy in your monitoring account. Additionally, the solution provides a **Service Quotas spoke template** and a **Trusted Advisor spoke template**. You must deploy each of these templates in the member accounts that need quota monitoring. The solution won't deploy the **Trusted Advisor spoke template** if the account doesn't have a support plan that includes the Trusted Advisor service.

Refer to Choose your deployment scenario for more information on how to use these templates, as well as two supplemental templates.

The **hub template** launches the following workflow:

1. **Reporting** - This workflow provisions an Amazon SNS topic, Amazon Simple Queue Service (Amazon SQS) queue, AWS Lambda function summarizer, and Amazon DynamoDB table. The queue receives usage events from all monitored accounts. The Lambda function puts all usage data on the DynamoDB table.
2. **Centralized event collection** - The workflow provisions a custom Amazon EventBridge bus, a corresponding rule, and Amazon SNS topic to raise alerts. The workflow raises alerts for quota usage and defines alert levels as:
   - **OK** (less than 80% utilization)
   - **WARN** (80% to 99% utilization)
   - **ERROR** (100% utilization)

   You can filter the alerts by excluding certain services or quotas through a notification configuration in AWS Systems Manager Parameter Store. The workflow also sends *all* events to the reporting queue for saving usage data in DynamoDB.
3. **Deployment management** - The workflow provisions Parameter Store, an Amazon EventBridge rule, a Lambda function, and CloudFormation StackSets. The workflow manages:
   - Permissions on the centralized EventBridge bus so that monitored accounts can send their usage events to it.
   - Deployment of spoke templates in the monitored accounts using StackSets when the solution is deployed in an organization (or OU).

   > **Note**
   >
   > When you update the Systems Manager parameter value with OU IDs or account IDs, the workflow makes needed configuration changes to start monitoring the updated list of OUs or accounts.

The **Service Quotas spoke template** launches the following workflow:

1. **Quota list generation** - The workflow provisions a Lambda function and two DynamoDB tables. The workflow manages an active and validated list of Service Quotas that support usage monitoring using CloudWatch metrics.
2. **Quota utilization alerting** - The workflow provisions a schedule-based Lambda function, custom EventBridge bus, and an Amazon EventBridge rule. The `CW Poller` function queries the quota list table and retrieves usage data for those quotas from CloudWatch metrics. The

workflow sends the usage data as events to the EventBridge bus. The spoke bus routes these usage events to both the centralized bus and the spoke SNS bus (if provided).

The **Spoke SNS template** launches the following workflow:

1. **Spoke account notifications** - The workflow provisions notification resources in the spoke account to decentralize notifications. Specifically, it provisions an EventBridge bus with a rule that routes messages to the SNS publisher Lambda function. This function applies any notification muting rules configured through the notification muting parameter in the SSM Parameter Store. The Lambda function then publishes the relevant events to the SNS topic in the spoke account.

The **Trusted Advisor spoke template** launches the following workflow:

1. **Trust Advisor alerting** - The workflow provisions a Lambda function and an Amazon EventBridge rule to support quota usage monitoring using Trusted Advisor. The Lambda function executes at an interval of 24 hours to refresh Trusted Advisor checks. The Events rule routes Trusted Advisor usage events to the centralized bus.

> **Note**
>
> AWS CloudFormation resources are created from AWS Cloud Development Kit (AWS CDK) components.

Architecture details

This section describes the components and AWS services that make up this solution and the architecture details on how these components work together.

AWS services in this solution

AWS service	Description
Amazon CloudWatch	**Core.** Monitors quota usage
Service Quotas	**Core.** Manages the quotas for your AWS services
AWS CloudFormation	**Core.** Deploys the solution templates in your account(s)
AWS Trusted Advisor	**Core.** Monitors quota usage and recommends resource deletion or quota increases
Amazon SNS	**Supporting.** Sends notification alerts when you reach the quota usage threshold
Amazon SQS	**Supporting.** Used as a dead-letter queue for asynchronously-invoked Lambda functions
AWS Lambda	**Supporting.** Deploys the functions to manage deployments, notificat ions, and querying quota usages
Amazon DynamoDB	**Supporting.** Deploys tables for the list of services, quotas monitored, and a summarizer
Amazon EventBridge	**Supporting.** Connects solution components by routing events
AWS Systems Manager	**Supporting.** Saves parameters such as notification configurations, OU IDs, or account IDs
AWS Organizations	**Optional.** Supports management of resources from manager and delegated administrator accounts

Plan your deployment

This section covers cost, security, quotas, AWS Regions, and other considerations for planning your deployment.

Supported AWS Regions

You can deploy the primary hub template (`quota-monitor-hub.template`), the Service Quotas spoke template (`quota-monitor-sq-spoke.template`), and supplemental prerequisite AWS CloudFormation templates in any AWS Region. You can deploy the Trusted Advisor template (`quota-monitor-ta-spoke.template`) only in the US East (N. Virginia) Region or the AWS GovCloud (US-West) Region.

Cost

The following tables provide a sample cost breakdown for deploying this solution with the default parameters in the US East (N. Virginia) Region for one month. There are additional minimal costs for a hub stack and a spoke stack.

> **Note**
>
> The monthly cost was estimated for a quota size of 2,000, as of this revision. If more quotas are supported by Service Quotas for additional resource types, the quota size would increase and monthly estimated cost could be higher. The costs provided are estimates based on a 6-hour monitoring frequency.
> The scalable costs (Amazon SQS and DynamoDB) increase with the number of accounts and Regions.

Monthly cost by number of accounts

Deployment size	Number of accounts	Number of AWS Regions	Cost per month [USD]
Small	10	8	$11.01 + 0.00355*10*8 + 0.01*10

Deployment size	Number of accounts	Number of AWS Regions	Cost per month [USD]
			+ $5.24*10*8 = **$430.39**
Medium	100	10	$11.01 + 0.00355*100*10 + $0.01*100 + 5.24*100*10 = **$5,253.55**
Large	1000	15	$11.01 + 0.00355*1000*15 + 0.01*1000 + 5.24*1000*15 = **$78,641.25**

Note

We calculated cost per month with the following formula: [monthly fixed cost for a hub stack] + [monthly scalable cost for a hub stack] * [number of accounts] * [number of regions] + [monthly cost for a Trusted Advisor spoke stack] * [number of accounts] + [monthly cost for a Service Quotas spoke stack] * [number of accounts] * [number of regions]

Monthly Fixed cost for a hub stack

AWS service	Cost per month [USD]
Amazon SNS topic	<$0.01
* AWS Lambda*	$10.00
* AWS KMS*	$1.00
Total cost:	**$11.01**

Monthly Scalable cost for a hub stack

AWS service	Cost per month [USD]
Amazon SQS queue	$0.00178
Amazon DynamoDB	$0.00177
Total cost:	**$0.00355**

Monthly cost for a Trusted Advisor spoke stack

AWS service	Cost per month [USD]
Amazon EventBridge	Free [1]
AWS Lambda	~ $0.01 [2]
Total cost:	**$ 0.01**

[1] AWS default service events are free. For more information, refer to Amazon EventBridge Pricing.

[2] The stack uses Support APIs which are not available under the free developer plan. For more information, refer to Compare Support Plans.

Monthly cost for a Service Quotas spoke stack

AWS service	Cost per month [USD]
Amazon EventBridge	$0.01
Amazon CloudWatch (`GetMetricData` API)	$5.12
AWS Lambda	~ $0.02
Amazon DynamoDB	~ $0.09
Total cost:	**$5.24**

Prices are subject to change. For full details, refer to the pricing webpage for each AWS service used in this solution.

> **Note**
>
> When you delete a stack, the DynamoDB table on the hub account is not deleted. DynamoDB will continue to incur costs until the you delete the table.

Security

When you build systems on AWS infrastructure, security responsibilities are shared between you and AWS. This shared responsibility model reduces your operational burden because AWS operates, manages, and controls the components including the host operating system, the virtualization layer, and the physical security of the facilities in which the services operate. For more information about security on AWS, visit AWS Cloud Security.

IAM roles

AWS Identity and Access Management (IAM) roles allow you to assign granular access policies and permissions to services and users on the AWS Cloud. This solution creates IAM roles that grant the solution's Lambda functions access to create Regional resources.

Quotas

The solution uses Trusted Advisor and Service Quotas to check quotas against resource utilization.

- **Trusted Advisor** - This solution supports 50 quota checks offered by Trusted Advisor. For more information, refer to Quota checks with Trusted Advisor.
- **Service Quotas** - This solution supports all quotas that allow resource usage monitoring using Amazon CloudWatch. When more quotas from different services start supporting resource usage monitoring, the solution automatically updates to support these new quotas. For details, refer to Service Quotas and Amazon CloudWatch alarms.

Slack integration

This solution includes an optional configuration to send notifications to your existing Slack channel. To use this feature, you must have an existing Slack channel and specify Slack webhook URL on the Systems Manager Parameter Store `/QuotaMonitor/SlackHook`.

The following figure depicts an example of using Slack notifications with the solution.

Image depicts an example Quota Monitor Notification in Slack

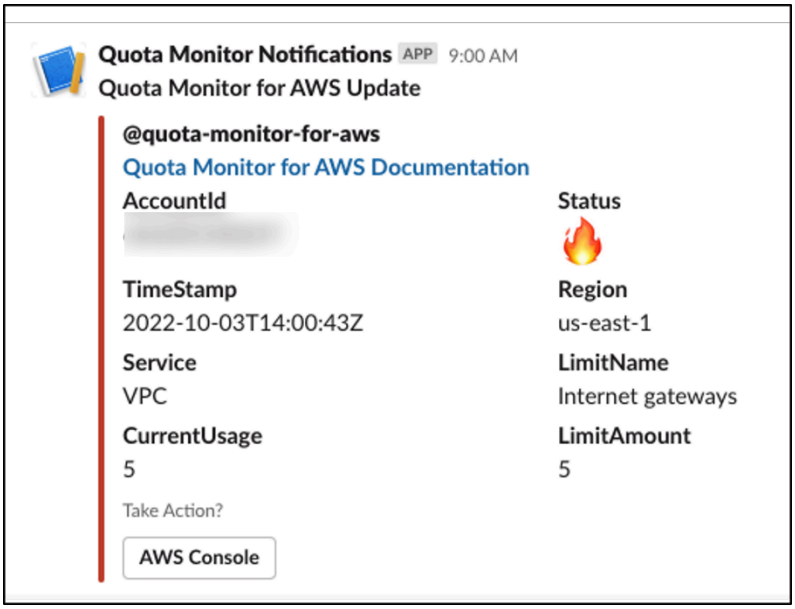

Amazon SQS dead-letter queue

The Quota Monitor for AWS solution deploys an Amazon SQS dead-letter queue. The `Summarizer` Lambda function, and other Lambda functions in the spoke accounts, attempt to process messages three times. If it cannot process the message after three attempts, it sends the message to the dead-letter queue where you can debug.

About Node.js versions

Quota Monitor for AWS version 5.3.0 and earlier versions use the Node.js 8.10 runtime, which reached end-of-life on December 31, 2019. Lambda now blocks both the create operation and the `update` operation. For more information, refer to Runtime Support Policy in the *AWS Lambda Developer Guide*. To continue using this solution with the latest features and improvements, update the stack as described in Update the solution.

Deployment scenarios

The solution supports different deployment scenarios for:

- Customers who use AWS Organizations
- Customers who don't use AWS Organizations
- Customers who just use individual AWS accounts
- Customers who use both AWS Organizations and individual AWS accounts.

For more information, refer to Choose your deployment scenario.

If you are deploying this solution in an environment with AWS Organizations, refer to Best practices for AWS Organizations

Spoke templates

The spoke templates packaged with the solution are standalone templates, and you can deploy them independently. To determine which templates to deploy, ask the following questions:

- Do you have a hub or monitoring account?
- Do you need the entire solution deployment with all of its features?

If you answered No to any of the questions above, then you can deploy just the spoke templates in the account to be monitored:

- `quota-monitor-ta-spoke.template` to support quota checks offered by Trusted Advisor
- `quota-monitor-sq-spoke.template` to support quota checks offered by Service Quotas

Additionally, the spoke templates offer extensions (such as sending notifications to different destinations). The spoke templates provision EventBridge rules for capturing OK, WARN, or ERROR quota events. You can configure these rules to send the events to destinations according to your requirements. For more details, refer to Amazon EventBridge targets.

Deploy the solution

This solution uses CloudFormation templates and stacks to automate its deployment. The CloudFormation templates specify the AWS resources included in this solution and their properties. The CloudFormation stack provisions the resources that are described in the templates.

Before you launch the automated deployment, review the architecture and other considerations discussed in this guide. Follow the step-by-step instructions in this section to configure and deploy the Quota Monitor for AWS into your account.

Time to deploy: Approximately five minutes.

Prerequisites

- If you are using AWS Organizations, you can deploy `quota-monitor-prerequisite.template` to fulfill needed prerequisites. For detailed instructions, refer to Step 2a: Launch the prerequisite stack (optional).
- To support quota usage monitoring with Trusted Advisor, each account must have a Business- or Enterprise-level Support plan to gain access to the Trusted Advisor service quota checks.
- To use this solution's Slack notification functionality, you must have an existing Slack channel.

> **Important**
>
> When deploying this solution across multiple account types (management, designated administrator for CloudFormation StackSets, and spoke accounts), ensure that the opted-in Regions overlap across all involved accounts. If the hub account has opt-in Regions that are not enabled in the spoke accounts, the solution will attempt to deploy in those Regions. This will cause deployment failures in the spoke accounts and might prevent successful deployment in the common Regions. Ensuring this Region overlap is crucial for the successful deployment and operation of the solution across your organization.

Deployment overview

To deploy this solution, follow these steps:

Step 1: Choose your deployment scenario

Choose the deployment scenario that meets your needs: AWS Organizations, hybrid, or AWS accounts that are not part of an AWS Organization.

Step 2a: Launch the prerequisite stack (optional)

Launch the prerequisite template in an Organizations management account to invoke a Lambda function that:

- Checks that the Organizations **All Features** is activated.
- Adds a member account as the designated administrator for CloudFormation StackSets.

--Or--

Step 2b: Fulfill the prerequisites manually (optional)

- Fulfill the prerequisites needed for monitoring quotas across Organizations manually.

Step 3a. Launch the hub stack for AWS Organizations

--Or--

Step 3b: Launch the hub stack for single account deployment

- Launch the AWS CloudFormation template into an AWS account that is registered as a delegated administrator for StackSets in your organization.
- Enter values for the required parameters: **Deployment Configuration.**
- Review the other template parameters and adjust if necessary.

Step 4a: Update the Systems Manager Parameter Store (Regions List)

- Update Parameter Store with the `RegionToDeploy`.

Step 4b: Update the Systems Manager Parameter Store (OUs)

- Update Parameter Store with the `OUs`.
- Review StackSets instances.

Step 5: Launch the spoke notifications stacks (optional)

- Launch the components necessary to add decentralized notification into each account.

Step 6: Launch the spoke stacks (optional)

- Launch the components necessary to monitor quotas in secondary accounts. Review the other template parameters and adjust if necessary.

Step 7: Configure notifications (optional)

- Configure notification filtering.

Step 8. Configure Slack notifications (optional)

- Configure Slack for notifications.
- Add the Slack webhook URL to the Systems Manager Parameter Store.

> **Important**
>
> This solution includes data collection. We use this data to better understand how customers use this solution and related services and products. AWS owns the data gathered though this survey. Data collection is subject to the AWS Privacy Notice.

AWS CloudFormation templates

This solution includes the following CloudFormation templates, which you can download before deployment:

View template

quota-monitor-hub.template - Use this template to launch the Quota Monitor for AWS solution and all associated components in the monitoring account.

View template

quota-monitor-sq-spoke.template - Use this template to launch the Quota Monitor for AWS solution and all associated components in secondary accounts to support Service Quotas.

View template

quota-monitor-sns-spoke.template - Use this template to launch notification resources in secondary accounts. This stack is optional and should be launched in only one Region within each secondary account.

View template

quota-monitor-ta-spoke.template - Use this template to launch the Quota Monitor for AWS solution and all associated components in secondary accounts to support Trusted Advisor.

View template

quota-monitor-prerequisite.template - Use this supplemental template to fulfill the prerequisites needed for monitoring quotas across AWS Organizations. This template should be launched in the organization management account.

View template

quota-monitor-hub-no-ou.template - Use this supplemental template to launch the Quota Monitor for AWS and all associated components in the monitoring account, when you are not using AWS Organizations.

Refer to Choose your deployment scenario later in this guide to determine which templates you need to deploy to meet your needs. Refer to the README.md file in the GitHub repository for guidance to customize the template.

Step 1. Choose your deployment scenario

You can deploy Quota Monitor for AWS in three deployment scenarios:

- **Scenario 1** - Environments where all AWS accounts are part of one or more Organizations.
- **Scenario 2** - Hybrid environments with Organizations and independent AWS accounts.
- **Scenario 3** - Environments that do not use Organizations and use single accounts instead.

To leverage all the benefits of this solution for automated monitoring and automated deployment, we recommend using Organizations.

The following sections describe how to deploy Quota Monitor for AWS in each of these deployment scenarios.

Deploying in AWS Organizations environments and hybrid environments (scenarios 1 and 2)

Choose this scenario if you are using Organizations and the AWS account that you are using for monitoring quotas is registered as a delegated administrator for StackSets in the organization.

You can choose from the two deployment modes provided as template input parameters:

- **Organizations (default mode)** - If you want to monitor quota utilization across your Organizations or across different OUs under your organization, choose this mode.
- **Hybrid** - If you want to monitor quota utilization across your Organizations, OUs, and accounts outside your Organizations, choose this mode.

The following figure depicts an example of deploying the solution in your monitoring account.

Image depicts the workflow for deploying a monitoring account

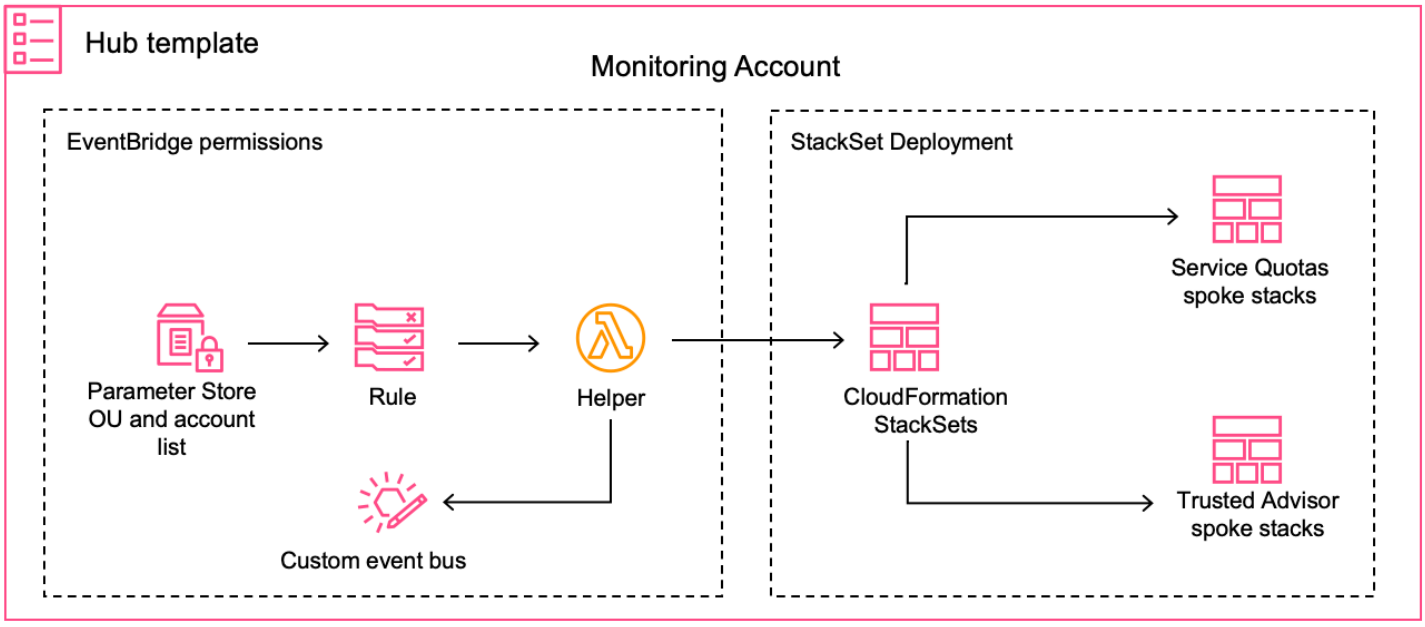

After you choose the deployment mode, the resources needed for that mode are provisioned. The deployment workflow is invoked when you update the deployed Systems Manager Parameter Store.

- The `helper` Lambda function updates the permissions on the centralized EventBridge bus, so all monitored accounts can send their quota utilization events to the monitoring account.
- CloudFormation StackSets automates spoke template deployments in the secondary accounts under targeted OUs.
- For additional accounts not under the purview of Organizations, you can manually deploy spoke templates.

Deploying when not using AWS Organizations (scenarios 2 and 3)

While we recommend using Organizations so that you can leverage the benefits of automated monitoring and automated deployment, you might have use cases where you are not using Organizations.

When you are not using Organizations and your monitoring account is not an organization member account—rather, an independent standalone account—use the supplemental `quota-monitor-hub-no-ou.template`.

> **Note**
>
> You are responsible for the cost of the AWS services used while running this solution. Review the Cost section for more details. For full details, refer to the pricing webpage for each AWS service you will be using in this solution.

The following flowchart depicts which templates you need to deploy, depending on your deployment scenario.

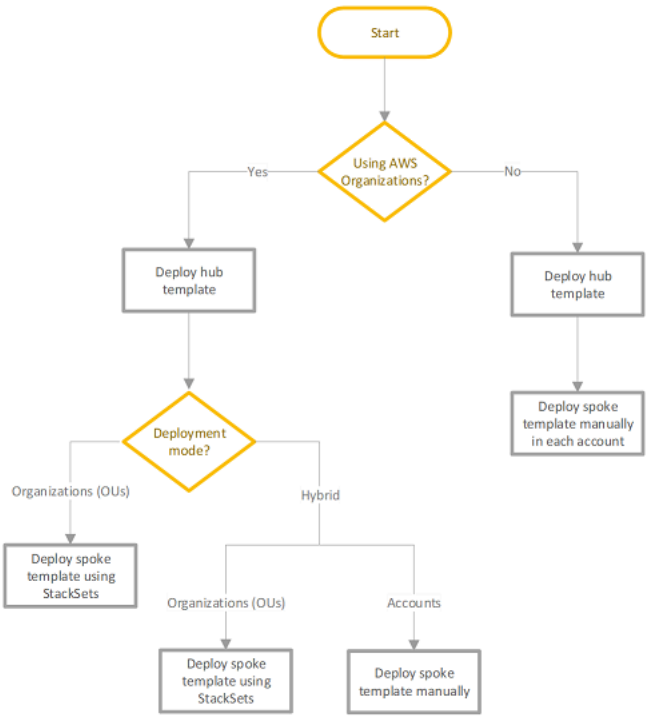

The following table summarizes the decision criteria for choosing templates, regions, and accounts for monitoring your quotas.

Question	Using AWS Organizations	Using single accounts
Where do you deploy a prerequisite template?	Deploy in a management account	N/A
Which AWS Region should you use for the prerequisite template?	Any AWS Region	N/A
Which hub template should you use?	`quota-monitor-hub.template`	`quota-monitor-hub-no-ou.template`
Which hub account should you use?	Any account	Any account
Where do you the deploy spoke templates?	* StackSets for Organizations and OU deployment scenarios	Manual deployment

Question	Using AWS Organizations	Using single accounts
	* StackSets and manual deployment for hybrid deployment scenarios	
Which spoke account should you use?	Any	Any
Which AWS Region should you use for the Trusted Advisor spoke template?	`us-east-1` or `us-gov-west-1` Region	`us-east-1` or `us-gov-west-1` Region
Which AWS Region should you use for the Service Quota spoke template?	Any	Any

Step 2a. Launch the prerequisite stack (optional)

> **Note**
>
> Use the prerequisite stack only for Organizations deployments.

The solution provides a supplemental prerequisite template. When you deploy this automated CloudFormation template in an Organizations management account, a Lambda function checks for the following prerequisites:

1. Checks that the **AWS Organizations All Features** is activated.
2. Adds a member account as the designated administrator for CloudFormation StackSets.

> **Note**
>
> The solution deploys service-managed StackSets. You must allow trusted access with AWS Organizations in the organization management account before you can use service-managed permissions on the AWS CloudFormation console (refer to Enable trusted access with AWS Organizations in the *AWS CloudFormation User Guide*) or AWS

Organizations console (refer to Enabling trusted access with AWS CloudFormation Stacksets in the *AWS Organizations User Guide*).

The Lambda function installs the prerequisites. If there are errors during prerequisite installation, a stack rollback occurs with an error message.

Use the following procedures to deploy the `quota-monitor-prerequisite.template` CloudFormation template.

1. Sign in to the AWS Management Console and select the button to launch the `quota-monitor-prerequisite.template` CloudFormation template.

> **Note**
>
> You must launch the template in the US East (N. Virginia) or AWS GovCloud (US-West) Region of the organization for the management account.

2. On the **Create stack** page, verify that the correct template URL is in the **Amazon S3 URL** text box and choose **Next**.
3. On the **Specify stack details** page, assign a name to your solution stack.
4. Under **Parameters**, review the parameters for this solution template and modify them as necessary. This solution uses the following default values.

Parameter	Default	Description
Quota Monitor Monitoring Account	*<Requires input>*	Account ID for the primary account. This account will also be configured as the StackSets administrator account.

1. Choose **Next**.

2. On the **Configure stack options*page, choose *Next**.
3. On the **Review** page, review and confirm the settings. Select the box acknowledging that the template will create IAM resources.
4. Choose **Create stack** to deploy the stack.

You can view the status of the stack in the AWS CloudFormation Console in the **Status** column. You should receive a `CREATE_COMPLETE` status in less than two minutes.

Step 2b. Fulfill prerequisites manually (optional)

> **Note**
>
> Use this procedure only for Organizations deployments.

Use the following procedure to manually fulfill prerequisites for the solution in your Organizations.

1. Activate **AWS Organizations Full Feature**.
2. Designate a member account as the StackSets administrator. This account will be your hub account.

> **Note**
>
> The solution deploys service-managed StackSets. You must allow trusted access with AWS Organizations in the organization management account before you can use service-managed permissions on the AWS CloudFormation console (refer to Enable trusted access with AWS Organizations in the *AWS CloudFormation User Guide*) or AWS Organizations console (refer to Enabling trusted access with AWS CloudFormation Stacksets in the *AWS Organizations User Guide*).

Step 3a. Launch the hub stack for AWS Organizations

This CloudFormation template deploys the Quota Monitor for AWS into your primary account.

> **Note**
>
> This template must be launched in a member account of your Organizations that is registered as delegated administrator for StackSets.
> You are responsible for the cost of the AWS services used while running this solution. Review the Cost section for more details. For full details, refer to the pricing webpage for each AWS service you will be using in this solution.

1. Sign in to the AWS Management Console and select the button to launch the `quota-monitor-hub.template` CloudFormation template.

   Launch solution

2. The template launches in the US East (N. Virginia) Region by default. To launch the solution in a different AWS Region, use the Region selector in the console navigation bar.

   > **Note**
   >
   > You can launch this template in any AWS Region.

3. On the **Create stack** page, verify that the correct template URL shows in the **Amazon S3 URL** text box and choose **Next**.
4. On the **Specify stack details** page, assign a name to your solution stack.
5. Under **Parameters**, review the parameters for the template and modify them as necessary. This solution uses the following default values.

Parameter	Default	Description
Deployment Configuration		
Deployment Mode	`Organizations`	The solution supports two deployment modes: **Organizations (default) - Monitor quota usage across your Organization.** Hybrid

Parameter	Default	Description
		- Monitor quota usage monitoring across your Organization and secondary accounts.
Regions List	`ALL`	The list of AWS Regions where the Service Quotas spoke templates should be deployed. Options are ALL or a comma separated list of AWS Regions (for example, `us-east-1`).
SNS Spoke Region	<Optional input>	Optionally, specify the AWS Region to launch the decentralized SNS stack in the spoke accounts. If left empty, the stack will not be launched.
Management Account Id	*	Organization's managemen t Id to scope permissions down for Stackset creation.
StackSet Deployment Options		
Region Concurrency	`PARALLEL`	Selection of whether to deploy the StackSets into AWS Regions in parallel (default) or sequentially.
Percentage Maximum concurrent accounts	`100`	Percentage of accounts per AWS Region to which you can deploy stacks at one time.

Parameter	Default	Description
Percentage Failure tolerance	`0`	Percentage of account, per AWS Region, for which stacks can fail before CloudFormation stops the operation in that Region.
Notification Configuration		
Email Notification	*`<Requires input>`*	Email address to receive alert notifications.
Slack Notification	No	Choose `Yes` if you want to receive Slack notifications for quota utilization alerts. You can specify the Slackhook web URL later in the Systems Manager Parameter Store (see Step 7: Configure notifications.
Stackset Stack Configuration Parameters		
Notification Threshold	`80`	The percentage threshold for quota utilization notifications.
Monitoring Frequency	`rate (12 hours)`	How often the quota monitoring scan should run. Available choices are every 6 or 12 hours.
Report OK Notifications	No	Whether to save the OK notifications in the summary table on the hub account.

Parameter	Default	Description
SageMaker AI Monitoring	Yes	Choose Yes to enable monitoring for SageMaker AI quotas, or No to disable it. Enabling SageMaker AI monitoring might consume a high number of quotas, which could lead to higher usage costs.
Connect Monitoring	Yes	Select Yes to enable monitoring for Amazon Connect quotas, or No to disable it. Keep in mind enabling Connect monitorin g might consume a high number of quotas, which could result in higher usage costs.

6. Choose **Next**.
7. On the **Configure stack options** page, choose **Next**.
8. On the **Review** page, review and confirm the settings. Select the box acknowledging that the template will create IAM resources.
9. Choose **Create stack** to deploy the stack.

   You can view the status of the stack in the AWS CloudFormation Console in the **Status** column. You should see a status of CREATE_COMPLETE in approximately five minutes.

10. In the subscription notification email, select the **SubscribeURL** link to turn on Amazon SNS email notifications.

- Customizing SageMaker and Amazon Connect services monitoring *

Because monitoring SageMaker and Amazon Connect services can incur high costs, this solution allows you to enable or disable monitoring for these services at the hub template level. This setting applies to all spoke accounts in your deployment.

To change these settings:

1. Update the hub stack in CloudFormation.
2. Modify the **SageMaker Monitoring** and **Connect Monitoring** parameters as needed.
3. Apply the stack update.

> **Note**
>
> Changing these parameters during a stack update affects all spoke accounts. If you leave them unchanged, the existing monitoring customizations in the spoke accounts remain intact.
> For spoke account-specific customization, you can modify the monitoring status in their Service DynamoDB table after deployment. The table includes entries for each service, such as SageMaker and Amazon Connect, with a **Monitored** field that can be set to `true` or `false`.

> **Important**
>
> Steps 4a and 4b are critical for the solution to function correctly. Without updating these, the solution won't know which accounts, OUs, or Regions to monitor.

Step 3b. Launch the hub stack for single account deployment

1. Sign in to the AWS Management Console and select the button to launch the `quota-monitor-hub-no-ou.template` CloudFormation template.

   Launch solution

2. The template launches in the US East (N. Virginia) Region by default. To launch the solution in a different AWS Region, use the Region selector in the console navigation bar.

   > **Note**
   >
   > You can launch this template in any AWS Region.

3. On the **Create stack** page, verify that the correct template URL shows in the **Amazon S3 URL** text box, and choose **Next**.
4. On the **Specify stack details** page, assign a name to your solution stack.
5. Under **Parameters**, review the parameter for the template and modify it as necessary. This solution uses the following default values.

Parameter	Default	Description
Email Notification	*<Optional input>*	Provide an email address to receive alert notifications.
Slack Notification	No	Choose `Yes` if you want to receive Slack notifications for quota utilization alerts.
Report OK Notifications	No	Whether to save the `OK` notifications in the summary table on the hub account.

1. Choose **Next**.
2. On the **Configure stack options** page, choose **Next**.
3. On the **Review** page, review and confirm the settings. Select the box acknowledging that the template will create IAM resources.
4. Choose **Create stack** to deploy the stack.

   You can view the status of the stack in the AWS CloudFormation Console in the **Status** column. You should see a status of CREATE_COMPLETE in approximately five minutes.
5. After the stack deploys, update the `/QuotaMonitor/Accounts` Systems Manager Parameter Store with list of targeted accounts. Once the parameter is updated:
   - The `helper` Lambda function updates the permissions on the centralized EventBridge bus so that all monitored accounts can send their quota utilization events to the primary account.
   - You can deploy `quota-monitor-ta-spoke` and `quota-monitor-sq-spoke` templates in the monitored accounts manually. Refer to Step 6: Launch the spoke stacks.

Step 4a. Update Systems Manager Parameter Store (Regions List)

Use the following procedure to update the Systems Manager Parameter Store with the list of AWS Regions where you want to deploy the spoke templates.

1. Open the AWS Systems Manager console.
2. In the navigation pane, choose **Parameter Store**.

Parameter Store

--Or--

If the Systems Manager home page opens first, choose the menu icon
(☰
) to open the navigation pane, then choose **Parameter Store**.

My Parameters

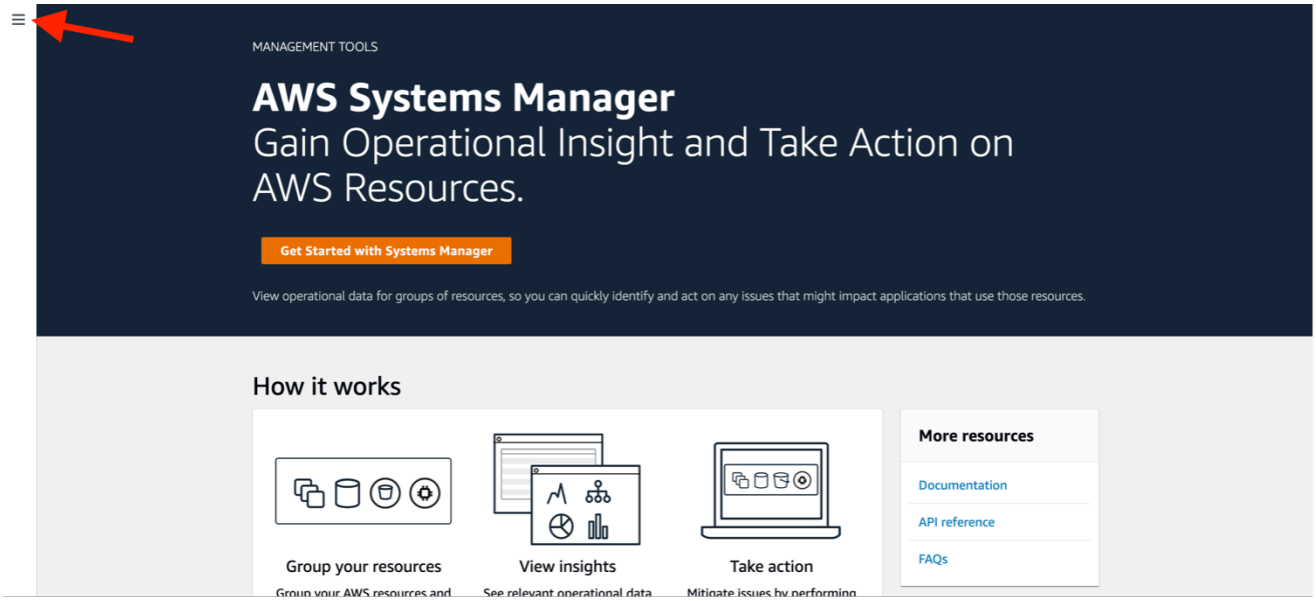

1. On the **My parameters** tab, select the box next to the parameter to update.
2. Choose **Edit**. Update the **Value**. The value should be comma-separated with no spaces. For example, `/QuotaMonitor/RegionsToDeploy: us-east-1,us-east-2`. The default value is `ALL`.
3. Choose **Save changes**.

Step 4b. Update Systems Manager Parameter Store (OUs)

Follow these steps to update the Systems Manager Parameter Store for the AWS accounts (**Account-Ids**) and OUs (**OU-ids**) you want to monitor.

1. Open the AWS Systems Manager console.
2. In the navigation pane, choose **Parameter Store**.

Parameter Store

--Or--

If the Systems Manager home page opens first, choose the menu icon (☰) to open the navigation pane, then choose **Parameter Store**.

My Parameters

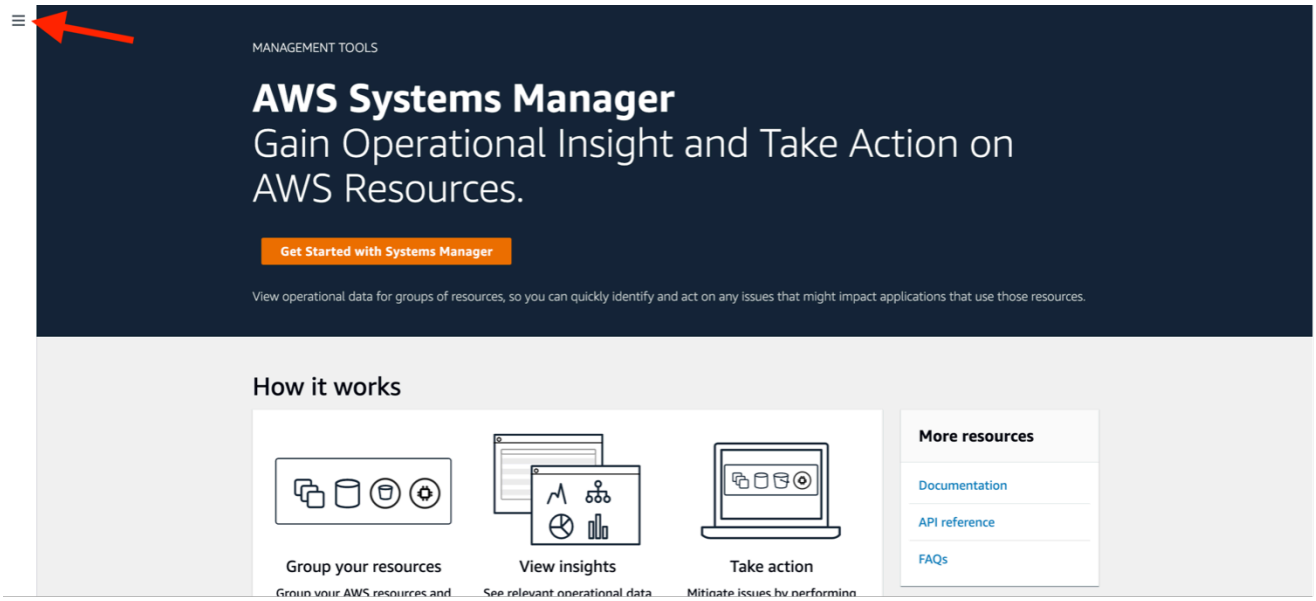

1. On the **My parameters** tab, select the box next to the parameter to update.

2. Choose **Edit**. Update the **Value**. The value should be comma-separated with no spaces For example, `/QuotaMonitor/OUs: ou-a1bc-d2efghij,ou-k1lm-n2opqrst`.
3. Choose **Save changes**.
4. Once you update the parameter, StackSets should start deploying solution templates in the targeted OUs or accounts. Review StackSets operation and instances.

Step 5. Launch the spoke SNS stacks when not using AWS Organizations or hybrid environments

Follow these steps to launch the components necessary for adding decentralized notifications to secondary accounts. Launch this stack in a single Region in every account where you want a separate SNS topic. Ensure that you launch this stack before the **sq spoke stacks** in Step 6.

> **Note**
>
> You are responsible for the cost of the AWS services used while running this solution. Review the Cost section for more details. For full details, refer to the pricing webpage for each AWS service you will be using in this solution.

1. Sign in to the AWS Management Console and select the button to launch the `quota-monitor-sns-spoke.template` CloudFormation template.

   Launch solution

2. The template launches in the US East (N. Virginia) Region by default. To launch the solution in a different AWS Region, use the Region selector in the console navigation bar.
3. On the **Create stack** page, verify that the correct template URL appears in the **Amazon S3 URL** text box, then choose **Next**.
4. On the **Specify stack details** page, assign a name to your solution stack.
5. Choose **Next**.
6. On the **Configure stack options** page, choose **Next**.
7. On the **Review** page, review and confirm the settings. Select the box acknowledging that the template will create IAM resources.

8. Choose **Create stack** to deploy the stack.

You can view the status of the stack in the AWS CloudFormation Console under the **Status** column. The status should show as `CREATE_COMPLETE` in approximately five minutes.

Step 6. Launch the spoke stacks when not using AWS Organizations or hybrid environments

Use the following procedure to launch the components necessary to monitor quotas in secondary accounts. You must launch the spoke stack in every account you want to monitor (including the account where the hub stack is deployed). You can deploy the Service Quotas spoke stack in all Regions, but only deploy the Trusted Advisor spoke stack in the Regions where the service's data plane resides, specifically, US East 1 (N. Virginia) and AWS GovCloud (US-West).

Enter the secondary account IDs in the `/QuotaMonitor/Accounts` Systems Manager Parameter Store provisioned by the primary template before you launch this template in secondary accounts. If you are using Organizations or a hybrid deployment mode, spoke template deployments are managed by CloudFormation StackSets.

> **Note**
>
> You are responsible for the cost of the AWS services used while running this solution. Review the Cost section for more details. For full details, refer to the pricing webpage for each AWS service you will be using in this solution.

1. Sign in to the AWS Management Console and select the button to launch the `quota-monitor-sq-spoke.template` CloudFormation template.

   Launch solution

2. The template is launched in the US East (N. Virginia) Region by default. To launch the solution in a different AWS Region, use the Region selector in the console navigation bar.

> **Note**
>
> You must launch `quota-monitor-ta-spoke.template` in the US East (N. Virginia) or AWS GovCloud (US-West) Region. You can launch the `quota-monitor-sq-spoke.template` in any AWS Region where you need quota monitoring.

3. On the **Create stack** page, verify that the correct template URL shows in the **Amazon S3 URL** text box, and choose **Next**.
4. On the **Specify stack details** page, assign a name to your solution stack.
5. Under **Parameters**, review the following parameter for the template and modify it as necessary.

Parameter	Default	Description
Event Bus Arn	*<Requires input>*	Amazon Resource Name (ARN) for the EventBridge bus in the primary account
Spoke SNS Region	<Optional input>	Optionally, specify the Region where you launched the spoke SNS stack in the spoke account. Ensure that the spoke SNS stack is launched in the spoke account first. Leave this parameter empty if you are not using the spoke SNS stack.

6. Choose **Next**.
7. On the **Configure stack options** page, choose **Next**.
8. On the **Review** page, review and confirm the settings. Select the box acknowledging that the template will create IAM resources.
9. Choose **Create stack** to deploy the stack.

   You can view the status of the stack in the AWS CloudFormation Console in the **Status** column. The status should show as `CREATE_COMPLETE` in approximately five minutes.

Step 7. Configure notifications (optional)

Follow these steps to configure and mute specific notifications for the Quota Monitor solution.

1. Open the AWS Systems Manager console.
2. In the navigation pane, choose **Parameter Store**.

   --Or--

   If the AWS Systems Manager home page opens first, choose the menu icon (☰) to open the navigation pane, and then choose **Parameter Store**.
3. On the **My parameters** tab, select the check box next to the parameter to update.
4. Choose **Edit**. Update the **Value**. The value should be comma-separated with no spaces The schema is `ServiceCode[:QuotaCode|QuotaName|Resource]`. Quotas matching that pattern will be muted such that no notification will be sent to the Amazon SNS topic or the Slack web hook. The following is an example:

   ```
   /QuotaMonitor/NotificationConfiguration: ec2:L-1216C47A,ec2:Running On-Demand
    Standard (A, C, D, H, I, M, R, T, Z) instances,dynamodb,logs:*,geo:L-05EFD12D.
   ```

   In this example, the following items occur:

   - The quotas `L-1216C47A` and `Running On-Demand Standard (A, C, D, H, I, M, R, T, Z) instances` from Amazon EC2 are muted.
   - All quotas from the DynamoDB are muted.
   - All quotas from the service logs are muted.
   - The quota `L-05EFD12D` from the service geo is muted.
5. Choose **Save changes**.

> **Note**
>
> You can get the values for the service code, quota code, quota name or resource from the notification email or slack message.

Muting Specific Notifications

If you receive notifications for quotas that are not useful or have very low limits, you can mute these notifications to avoid unnecessary alerts. As an example, this guide will walk you through the process of muting notifications for the `StartAutomationExecution` API in SSM, which has a limit of 1, using the Quota Monitor solution.

Identify the Quota to Mute: The quota in question is Transactions per second (TPS) for the `StartAutomationExecution` API with the limit code `L-99469188`.

Update the Notification Configuration: Follow the steps above to edit the **/QuotaMonitor/NotificationConfiguration** parameter and add `SSM:L-99469188` to the list.

Identifying Quotas with Limit 1

Service quotas in AWS set a limit of 1 for certain resources to provide highly available and reliable service to all customers. These limits are designed to protect customers from unintentional spend and excessive provisioning. There are two ways to identify which quotas have a limit of 1:

1. Check the ServiceQuotas Table: Go to the `SQQuotaTable` DynamoDB table and see the `Value` column for quotas, then sort it.

2. Run a Script: Run this script from our Github repo to get the list of quotas which have a limit of 1.

Step 8. Configure Slack notifications (optional)

1. Navigate to your workspace's Slack app.

   If required, sign in to Slack.
2. Choose **Create New App**.
3. Choose **From Scratch**.
4. Give the app a name and assign it to your workspace.
5. In the **Add features and functionality** section, select **Incoming Webhooks**.
6. Allow the feature and choose **Add New Webhook to Workspace**.
7. In the **Post to Channel** dropdown menu, select a channel.
8. Copy the WebHook URL.

9. In the AWS Systems Manager console, under **Shared Resources** in the left pane, select **Parameter Store**.
10. Select the `/QuotaMonitor/SlackHook` parameter, then choose **Edit**.
11. Update the value with your WebHook URL and choose **Save changes**.

Global China Region (GCR) deployment

You can deploy the Quota Monitor for AWS solution in AWS China Regions (Beijing and Ningxia) with certain regional limitations and considerations.

Limitations in China Regions:

- EventBridge does not support cross-region event routing.
- The Trusted Advisor (TA) stack is not supported in China Regions.

Deployment strategy for China Regions

To accommodate these limitations, follow this deployment strategy for the hybrid/OU model:

1. Hub deployment:
   - Deploy the hub stack separately in `cn-north-1` (Beijing) and `cn-northwest-1` (Ningxia) if you want to monitor services in both Regions.
   - Use the `quota-monitor-hub.template` CloudFormation template for the hub deployment.
2. Spoke deployment:
   - Deploy spoke stacks in the same Region as their corresponding hub.
   - Use the `quota-monitor-spoke.template` CloudFormation template for spoke deployment.

Important

- All deployments are Region-specific and do not support cross-region monitoring.
- The hub and associated spoke stacks must be deployed in the same Region.
- To monitor supported services in both China Regions, deploy the solution twice - once in each Region.

Deployment models

1. Account model:
   - Use the `quota-monitor-hub-no-ou.template` CloudFormation template for single account deployments.
   - Use this model when deploying Quota Monitor for individual accounts.
   - Deploy spoke stacks manually in the same Region as the hub using the GCR-specific spoke template.
   - Follow the steps in the Deploy the solution section for more information.
2. Hybrid/OU model:
   - Use this model when deploying across an AWS Organization or a mix of Organization and individual accounts.
   - In the CloudFormation template, specify the Region where you're deploying the hub in the Region parameter.
   - If you leave the default value `ALL` for the Regions parameter, the solution will attempt to deploy StackSets in both China Regions. The deployment will succeed in the current hub Region but fail in the other Region. The solution will still function correctly, monitoring services in the current hub Region for all spoke stacks.

> **Note**
>
> All monitored accounts must be in the same China Region as the hub.

For detailed steps on deploying hub and spoke stacks, refer to the Deploy the solution section. Follow those steps for each China Region where you want to deploy, using the GCR-specific templates provided above.

> **Note**
>
> Some features available in global Regions might not be supported in China Regions. Always refer to the AWS documentation for the most up-to-date information on service availability in China Regions.

Update the solution

If you previously deployed the solution, follow these procedures to get the latest version of the solution's framework.

Update to v6.1.0 or later

Starting from version 6.1.0 of the solution, we used AWS managed keys on the spoke stacks, which previously used customer managed keys. The hub stacks still use customer managed keys. This helps your cost because AWS managed keys don't incur the monthly storage fee that customer managed keys incur.

Updating the stack and the StackSet instances, or deleting them, won't delete the customer managed keys. You must manually delete the following customer managed keys after updating the solution from a version that uses customer managed keys, or after deleting the solution:

- One customer managed key for the hub stack.
- One customer managed key per account, per AWS Region, for the spoke stack.

Update the hub stack

Use the following procedure to update your main AWS CloudFormation stack to the current version.

1. From your main account where the Quota Monitor for AWS hub template deployed, sign in to the AWS CloudFormation console.
2. From the **Stacks** page, select your Quota Monitor for AWS hub stack and choose **Update**.

> **Note**
>
> Take a note of the parameter **Account List**. You'll need this later to complete the update process.

3. On the **Update stack** page, verify that **Replace current template** is selected.
   - In the **Specify template section, select *Amazon S3 URL**.
   - Copy the link of the latest template for `quota-monitor-hub.template`.

- Paste the link in the **Amazon S3 URL** box.
- Verify that the correct template URL shows in the **Amazon S3 URL** box, and choose **Next**.

4. On the **Specify stack details** page, under **Parameters**, review the parameters for the template and modify them as necessary. Refer to Step 4a. Update the Systems Manager Parameter Store (Regions list) for details about the parameters.

> **Note**
>
> When updating the SNS Spoke Region parameter:

- Specify only one Region. If you provide multiple Regions, StackSet instances won't deploy, and you'll see warnings in the deployment manager logs.
- Changing this parameter doesn't initiate an immediate deployment or redeployment of the SNS spoke stack.

5. Choose **Next**.
6. On the **Configure stack options** page, choose **Next**.
7. On the **Review** page, review and confirm the settings. Select the box acknowledging that the template will create IAM resources.
8. Choose **View change set** and verify the changes.
9. Choose **Update stack** to deploy the stack.

You can view the status of the stack in the AWS CloudFormation console in the **Status** column. You should see a status of **UPDATE_COMPLETE** in approximately five minutes.

Once the stack status is **UPDATE_COMPLETE**, follow these additional steps:

1. Update the `/QuotaMonitor/Accounts` Systems Manager Parameter Store with the **Account List** value copied earlier.
2. If you opted-in for Amazon SNS email notifications, you should receive email to confirm subscription.

After updating the hub stack, especially when changing the SNS Spoke Region, follow these steps:

1. Open the AWS Systems Manager console.

2. Go to Parameter Store.
3. Locate and update (or resave without making changes) the parameter for OUs (/QuotaMonitor/OUs) or Accounts (/QuotaMonitor/Accounts), depending on your deployment model. This action triggers the deployment manager to:
   - Deploy the SNS spoke stack in the new Region.
   - Remove the SNS spoke stack from the old Region (if applicable).

Impact on SageMaker AI and Amazon Connect monitoring during stack updates

In the Hybrid/OU model, when you update the hub stack, changes to the **SageMaker AI Monitoring** and **Connect Monitoring** parameters affect your deployment as follows:

1. If you modify the SageMaker AI Monitoring or Connect Monitoring parameters during the update, the new settings will apply to all spoke accounts.
2. If you leave these parameters unchanged during the update:
   - Existing spoke accounts will retain their current SageMaker AI and Amazon Connect monitoring settings.
   - New spoke accounts added after the update will inherit the current stack-level monitoring settings.
3. If you change the **SageMaker AI Monitoring** and **Connect Monitoring** parameters during the update, any spoke account-specific monitoring changes made for the two services in the Service DynamoDB table will be overwritten.

To preserve account-specific monitoring settings while updating other aspects of the stack, leave the SageMaker AI Monitoring and Connect Monitoring parameters unchanged during the update.

Update the spoke stack if you aren't using Organizations

The solution previously did not support Organizations. If you are continuing with that model where your hub account is an independent account and not part of an organization, then follow instructions in this section. If you decide to use Organizations, then follow the launch steps in Step 1. Choose your deployment scenario.

This section provides instructions for updating the hub and spoke CloudFormation stacks if you don't use Organizations.

> **Note**
>
> When deleting a stack, your information in DynamoDB will not be deleted, and you will continue to incur costs until you delete the DynamoDB tables.

When you deploy the solution without Organizations, use the following procedure to update the spoke CloudFormation stack in your secondary accounts to the current version of `quota-monitor-ta-spoke.template` or `quota-monitor-sq-spoke.template`.

1. From your secondary account where the Quota Monitor for AWS spoke template is deployed, sign in to the CloudFormation console.
2. From the **Stacks** page, select your primary Quota Monitor for AWS stack and choose **Update**.
3. On the **Update stack** page, verify that **Replace current template** is selected.
4. In the **Specify template** section, select **Amazon S3 URL**.
5. Copy the link of the latest template for `quota-monitor-ta-spoke.template` or `quota-monitor-sq-spoke.template`.
6. Paste the link in the **Amazon S3 URL** box.
7. Verify that the correct template URL shows in the **Amazon S3 URL** box and choose **Next**.
8. On the **Specify stack details** page, under **Parameters**, review the parameters for the template and modify them as necessary. Refer to Step 6. Launch the spoke stacks for details about the parameters.

   > **Note**
   >
   > You can copy the EventBridge bus ARN from the **Outputs** section of the hub stack.

9. Choose **Next**.
10. On the **Configure stack options** page, choose **Next**.
11. On the **Review** page, review and confirm the settings. Select the box acknowledging that the template will create IAM resources.
12. Choose **View change set** and verify the changes.
13. Choose **Update stack** to deploy the stack.

Repeat these steps to update additional secondary accounts containing the spoke template.

Update or edit the spoke StackSets if you are using Organizations

When you deploy the solution in an organization, edit the StackSet details to update the spoke stacks in one operation. With the v6.0.0 release, there are two spoke templates; previous release had one spoke template. The two spoke templates individually own and provision resources needed to support Trusted Advisor and Service Quotas checks.

Use the following procedure to update the spoke AWS CloudFormation stack in your secondary accounts to the current version of `quota-monitor-ta-spoke.template` or `quota-monitor-sq-spoke.template`.

1. From your monitoring account where the Quota Monitor for AWS spoke template is deployed, sign in to the CloudFormation console.
2. From the **StackSets** page, select your primary Quota Monitor for AWS stack and choose **Edit StackSet details**.
3. On the **Update stack** page, verify that **Replace current template** is selected.
   - In the **Specify template** section, select **Amazon S3 URL**.
   - Copy the link of the latest template for `quota-monitor-ta-spoke.template` or `quota-monitor-sq-spoke.template`.
   - Paste the link in the **Amazon S3 URL** box.
   - Verify that the correct template URL shows in the **Amazon S3 URL** box, and choose **Next**.
4. On the **Specify stack details** page, under **Parameters**, review the parameters for the template and modify them as necessary.

   > **Note**
   >
   > You can copy the EventBridge bus ARN from the **Outputs** section of the hub stack.

5. Choose **Next**.
6. On the **Configure stack options** page, choose **Next**.
7. On the **Set deployment options** page:
   - Specify organization unit ID(s)

- Specify AWS Regions
- Specify deployment options

8. On the **Review** page, review and confirm the settings. Select the box acknowledging that the template will create IAM resources.
9. Choose **Submit** to deploy the update.

You can update stacks in secondary accounts directly within those accounts.

Troubleshooting

The Quota Monitor for AWS logs errors, warnings, informational messages, and debugging messages for the solution's Lambda functions. To choose the type of messages to log, find the applicable function in the Lambda console, and change the **LOG_LEVEL** environment variable to the applicable type of message.

Level	Description
ERROR	Logs will include information on anything that causes an operation to fail.
WARNING	Logs will include information on anything that can potentially cause inconsistencies in the function but might not cause the operation to fail. Logs will also include ERROR messages.
NFO	Logs will include high-level information about how the function is operating. Logs will also include `ERROR` and `WARNING` messages.
DEBUG	Logs will include information that might be helpful when debugging a problem with the function. Logs will also include `ERROR`, `WARNING`, and `INFO` messages.

If these instructions don't address your issue, see the Contact AWS Support section for instructions on opening an AWS Support case for this solution.

Problem: Solution is not monitoring expected accounts or Regions

If you've deployed the solution but it's not monitoring the accounts or Regions, and the stacks are not deployed in spoke accounts, follow these steps to troubleshoot:

1. SSM parameter updates: Verify that you've updated the following SSM parameters after deployment:
   - `/QuotaMonitor/OUs` (for Organizations or Hybrid deployment)
   - `/QuotaMonitor/Accounts` (for Account or Hybrid deployment)
   - `/QuotaMonitor/RegionsToDeploy`
2. Parameter values: Ensure the values in these parameters are correct and formatted properly (comma-separated lists).
3. Deployment model: Confirm that you've selected the correct deployment model (Organizations, Account, or Hybrid) when creating the stack.
4. StackSet deployments: If you're using Organizations or Hybrid mode, check the CloudFormation StackSets console to confirm that the stacks have been created in the expected accounts and Regions.

Resolution

If any of the above are incorrect, update the SSM parameters with the correct values. The solution should detect these changes and adjust its deployment accordingly. If issues persist, update the main stack to trigger a redeployment.

Problem: The SNS spoke stack is not deploying in the new Region after updating the hub stack

After updating the hub stack, if the SNS spoke stack is not deploying in the new Region, follow these steps to troubleshoot:

Resolution

Ensure that you provide only one Region and update (or resave) the OUs or account's parameter in Systems Manager Parameter Store after changing the SNS Spoke Region in the hub stack.

Problem: Amazon CloudWatch Events bus permissions error

If during spoke stack deployment, you received a **CREATE_FAILED** message for the `TAWarnRule` and/or the `TASErrorRule`, verify that the CloudWatch Event Bus in the primary account allows the spoke account to send events to the monitoring account.

Resolution

Update the hub stack with the secondary account ID or complete the following tasks:

1. In the monitoring account, navigate to the Amazon CloudWatch console.
2. In the navigation pane, select **Event Buses**.
3. Select **Add Permissions**.
4. For **Principal**, enter the applicable secondary account ID.
5. Select the **Everybody()** box.
6. Choose **Add**.

Problem: Slack notifications are not being received

If you don't receive Slack notifications for WARN or ERROR events, check the CloudWatch logs for an error message.

1. In the primary account, navigate to the Amazon CloudWatch console.
2. In the navigation pane, select **Logs**.
3. Select the `/aws/lambda/[replaceable]<stackname>`-SlackNotifier-`[replaceable]<randomstring>` Log Group.
4. Select the top (most recent) Log Stream.
5. Look for the following error.

Example error for Quota Monitor for AWS

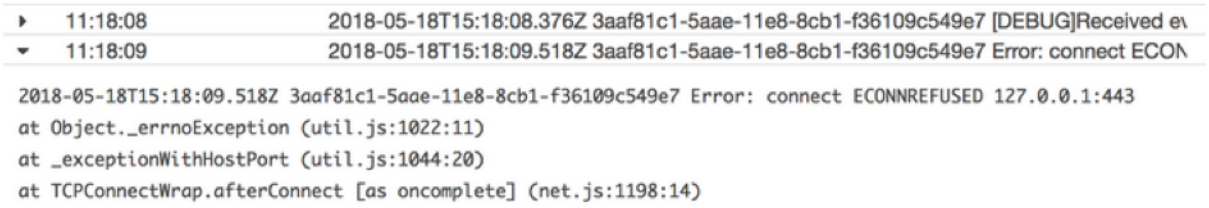

Resolution

Complete the following tasks:

1. In the primary account, navigate to the AWS Systems Manager console.

2. In the navigation pane under **Shared Resources**, select **Parameter Store**.
3. Select the `/QuotaMonitor/SlackHook` parameter and verify that the parameter shows the correct value.

Problem: Email notifications are not being received

If you don't receive email notifications, confirm that you subscribed to the Amazon SNS topic.

1. In the primary account, navigate to the Amazon SNS console.
2. In the navigation pane, select **Topics**.
3. Select the `<stackname>-SNSTopic-<randomstring>` ARN value.
4. Verify that the **Subscription ID** shows an ARN value.

Resolution

If the **Subscription ID** field shows `PendingConfirmation`, complete the following tasks:

1. Select the box next to **PendingConfirmation**.
2. Under **Subscriptions**, select **Request Confirmations**.
3. Navigate to the applicable email inbox.
4. In the subscription notification email, select the **SubscribeURL** link.
5. In the Amazon SNS console, refresh and verify that the **Subscription ID** has an ARN value.

Problem: Hub stack creation failed

If the hub stack creation failed with the following error, you haven't allowed trusted access with Organizations:

```
You must enable organizations access to operate a service managed stack set
(Service: CloudFormation, Status Code: 400, Request ID: ABCXYZ
```

Resolution

Allow trusted access with AWS Organizations to use service-managed permissions on AWS CloudFormation console or AWS Organizations console. See Step 2b: Fulfill prerequisites manually for instructions.

Problem: Too many messages queued in the Summarizer SQS queue

If too many messages are queued in the QMSummarizerEventQueueQMSummarizerEventQueue SQS queue and the number of queued messages keeps growing.

Resolution

The QMReporterQMReporterLambda Lambda function consumes events from the queue and is invoked every five minutes by default. Try one or more of the following:

- Increase the rate of the QMReporterQMReporterEvents EventBridge rule on the default bus.
- Increase the value of the Lambda function's MAX_LOOPS environmental variable.

Problem: Identify the hub account number that centralizes quota alerts for the spoke account

In Organization/Hybrid Mode, it can be challenging to determine which hub account is responsible for centralizing quota alerts for a specific spoke account.

Resolution

To identify the hub account number:

1. Log in to the spoke account.
2. Navigate to the Amazon EventBridge console.
3. Select the 'QuotaMonitorSpokeBus' event bus.
4. Look for a rule named 'QMUtilizationWarn' or 'QMUtilizationErr'.
5. Click on the identified rule.
6. Examine the target event bus ARN in the rule details.

The account number in the target event bus ARN is the hub account number.

Example ARN

```
arn:aws:events:us-east-1:123456789012:event-bus/QuotaMonitorBus
```

In this example, 123456789012 is the hub account number.

Problem: Stack failed to update due to deletion of resources outside of CloudFormation

If the solution's resources are manually deleted outside of the CloudFormation stack, the solution will fail to update because the CloudFormation stack will not be able to find the resource.

Resolution

For resolution, see the How do I update a CloudFormation stack that's failing because of a resource that I manually deleted? article on AWS re:Post.

Contact AWS Support

If you have AWS Business Support+, AWS Enterprise Support, or AWS Unified Operations, you can use AWS Support Center to get expert assistance with this solution. The following sections provide instructions.

Create case

1. Sign in to Support Center.
2. Choose **Create case**.

How can we help?

1. Choose **Technical**.
2. For **Service**, select **Solutions**.
3. For **Category**, select **Other Solutions**.
4. For **Severity**, select the option that best matches your use case.
5. When you enter the **Service**, **Category**, and **Severity**, the interface populates links to common troubleshooting questions. If you can't resolve your question with these links, choose **Next step: Additional information**.

Additional information

1. For **Subject**, enter text summarizing your question or issue.
2. For **Description**, describe the issue in detail, including the name of this solution: Quota Monitor for AWS v6.3.10.
3. Choose **Attach files**.
4. Attach the information that Support needs to process the request.

Help us resolve your case faster

1. Enter the requested information.
2. Choose **Next step: Solve now or contact us**.

Solve now or contact us

1. Review the **Solve now** solutions.
2. If you can't resolve your issue with these solutions, choose **Contact us**, enter the requested information, and choose **Submit**.

Uninstall the solution

You can uninstall the Quota Monitor for AWS solution from the AWS Management Console or by using the AWS Command Line Interface (AWS CLI).

Note

You must manually delete the following:

- Delete the DynamoDB summary table from the hub account. The solution doesn't automatically delete the table, allowing you to maintain historical analysis of the quota usage and alert notifications if desired.
- Delete the customer managed key for the hub stack.
- Delete one customer managed key per account, per AWS Region, for the spoke stack.
- If you deployed the solution with Organizations, delete the StackSet instances from the StackSets before you delete the hub stacks.

Using the AWS Management Console

1. Sign in to the AWS CloudFormation console.
2. On the **Stacks** page, select this solution's installation stack.
3. Choose **Delete**.

Using AWS Command Line Interface

Determine whether the AWS CLI is available in your environment. For installation instructions, see What Is the AWS Command Line Interface in the *AWS CLI User Guide*. After confirming that the AWS CLI is available, run the following command.

```
$ aws cloudformation delete-stack --stack-name <installation-stack-name>
```

Deleting StackSet instances

You can delete the StackSet instances from the AWS Management Console or by using the AWS CLI.

Using the AWS Management Console

1. Sign in to the AWS CloudFormation console.
2. On the **StackSets** page, select this solution's installation StackSet.
3. Choose **Actions**, then choose **Delete stacks from StackSet**.

Using AWS Command Line Interface

Determine whether the AWS CLI is available in your environment. For installation instructions, see What Is the AWS Command Line Interface in the *AWS CLI User Guide*. After confirming that the AWS CLI is available, run the following command.

```
$ aws cloudformation delete-stack-instances -stack-set-name <installation-stackset-name> --regions <value>
```

Deleting the DynamoDB table

This solution is configured to retain the solution-created DynamoDB table. Follow these steps to delete the DynamoDB table.

1. Sign in to the DynamoDB console.
2. Choose **Tables** from the left navigation pane.
3. Locate the *<stack-name>* prefixed table and choose **Delete**.

To delete the DynamoDB table using AWS CLI, run the following command:

```
$ aws dynamodb delete-table <table-name>
```

Deleting the customer managed keys (scheduling deletion)

This solution is configured to retain the solution-created customer managed keys along with the DynamoDB table. Follow these steps to delete the customer managed keys.

1. Sign in to the AWS KMS console.
2. Choose **Customer managed keys** from the left navigation pane.
3. Locate the *<CMK-stack-name>* prefixed table and choose ***ey Actions and schedule key deletion**.

To delete the customer managed keys using AWS CLI, run the following command:

```
$ aws kms schedule-key-deletion --key-id <key-id>
```

Developer guide

This section provides the source code for the solution and additional customizations.

Source code

Visit our GitHub repository to download the source files for this solution and to share your customizations with others. The Quota Monitor for AWS templates are generated using the AWS CDK. Refer to the README.md file for more information.

Understanding Solution Components

Monitored Services and Quotas

> **Note**
>
> The solution only monitors quotas that have utilization metrics available through the AWS Service Quotas API.

The solution uses two main DynamoDB tables to manage services and quotas:

SQServiceTable

This table contains all AWS services that support Service Quotas. Note that a service being in this table doesn't necessarily mean it has quotas with utilization metrics. The **Monitored** flag indicates if the solution is actively tracking quotas for a service.

SQQuotaTable

This table contains the quotas actively monitored by the solution. You can filter by services or quotas to view detailed monitoring information. Note that if a service in SQServiceTable is marked as 'Monitored' but has no corresponding quotas in SQQuotaTable, it indicates that the service lacks quotas supporting utilization metrics and is therefore not monitored by the solution.

To modify monitoring:

1. Login to spoke acocunt and in **SQServiceTable**, find the service and set **Monitored** to true/false.

2. The solution automatically updates **SQQuotaTable** accordingly.

Summary DynamoDB Table

A table in the hub account (named 'QM-Table') stores quota utilization events.

Key features:

- Collects data from all monitored accounts and regions
- Fields: AccountId, Region, Service, Resource, LimitCode, LimitName, CurrentUsage, LimitAmount, Status
- TTL: 15 days

This table allows for historical analysis of quota usage across the organization.

Change the Lambda interval

By default, a Lambda function runs once every 24 hours to refresh the Trusted Advisor service quota checks. You can change how often the Lambda function is invoked by changing the mapping in the `quota-monitor-ta-spoke.template` CloudFormation template.

> **Note**
>
> If you set the Lambda function to run more than once every 12 hours, you may experience a delay in your events arriving in DynamoDB.

To change the interval, modify the following mapping in the CloudFormation template:

```
RefreshRate:
    CronSchedule:
      Default: rate(1 day)
```

For more information about acceptable rate expressions, refer to Rate Expressions in the *Amazon EventBridge User Guide.*

The same value is exposed as a template parameter for the `quota-monitor-sq-spoke.template` CloudFormation template.

Testing notifications using test events

This solution allows you to send test events using Lambda test events in the same way that a Service Quotas event would be captured. Follow these steps to send a test event:

1. Navigate to the Lambda Console in the AWS account you want to test.
2. Locate and select the `CW Poller` Lambda function.
3. Choose the **Test** tab.
4. In the **Event JSON** window, input the following event (set test-type to either `ERROR` or `WARN`):

```
{
  "detail-type": "QM Lambda Test Event",
  "test-type": "WARN"
}
```

5. Choose **Test**.

The event should send a notification to subscribers of the SNS topic.

Reference

This section includes information about data collection and a list of builders who contributed to this solution.

Data collection

This solution sends operational metrics to AWS (the "Data") about the use of this solution. We use this Data to better understand how customers use this solution and related services and products. AWS's collection of this Data is subject to the AWS Privacy Notice.

Contributors

The following individuals contributed to this document:

- Aaron Schuetter
- Garvit Singh
- Abe Wubshet
- George Lenz
- Sanjay Reddy Kandi

Revisions

Visit the CHANGELOG.md in our GitHub repository to track version-specific improvements and fixes.

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents AWS current product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided "as is" without warranties, representations, or conditions of any kind, whether express or implied. AWS responsibilities and liabilities to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Quota Monitor for AWS is licensed under the terms of the Apache License Version 2.0.