

User Guide

Amazon Managed Service for Prometheus



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

Amazon Managed Service for Prometheus: User Guide

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

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

Table of Contents

What is Amazon Managed Service for Prometheus?	1
Supported Regions	1
Pricing	8
Premium support	8
Get started	9
Set up AWS	9
Sign up for an AWS account	. 10
Create a user with administrative access	. 10
Create a workspace	. 11
Ingest metrics	. 12
Step 1: Add new Helm chart repositories	13
Step 2: Create a Prometheus namespace	. 14
Step 3: Set up IAM roles for service accounts	. 14
Step 4: Set up the new server and start ingesting metrics	. 14
Query metrics	16
Manage workspaces	. 18
Create a workspace	. 18
Configure your workspace	. 21
Edit a workspace alias	. 22
Find your workspace details	23
Delete a workspace	. 25
Ingest metrics	26
AWS managed collectors	. 27
Using a managed collector	. 28
Prometheus-compatible metrics	. 47
Customer managed collectors	. 48
Secure the ingestion of your metrics	48
ADOT collectors	. 49
Prometheus collectors	66
High-availability data	. 75
Query your metrics	. 83
Secure your metric queries	83
Using AWS PrivateLink with Amazon Managed Service for Prometheus	49
Authentication and authorization	49

Use Amazon Managed Grafana	
Connecting to Amazon Managed Grafana in a private VPC	85
Use Grafana open source	
Prerequisites	85
Step 1: Set up AWS SigV4	86
Step 2: Add the Prometheus data source in Grafana	87
Step 3: (optional) Troubleshooting if Save & Test doesn't work	
Use Grafana in Amazon EKS	90
Set up AWS SigV4	90
Set up IAM roles for service accounts	91
Upgrade the Grafana server using Helm	92
Add the Prometheus data source in Grafana	
Use direct queries	94
Query with awscurl	
Query statistics	97
Recording and alerting rules	100
Necessary IAM permissions	101
Create a rules file	102
Upload a rules file	104
Edit a rules file	106
Troubleshooting Ruler	107
Alert manager	109
Necessary IAM permissions	110
Create a configuration file	111
Set up an alert receiver	113
Create Amazon SNS topic	114
Amazon SNS permissions needed	114
Send alerts to your Amazon SNS topic	117
Send messages as JSON	119
Send alerts to other destinations	120
Amazon SNS validation rules	122
Upload a configuration file	123
Integrate alerts with Grafana	126
Prerequisites	126
Setting up Amazon Managed Grafana	127
Troubleshoot alert manager	128

Empty content warning	128
Non ASCII warning	129
Invalid key/value warning	129
Message limit warning	130
No resource based policy error	130
Not authorized to call KMS	131
Monitoring workspaces	132
CloudWatch metrics	132
Setting a CloudWatch alarm	139
CloudWatch Logs	139
Configuring CloudWatch Logs	140
Query insights and control	142
Configuring query logging	143
Configuring query throttling thresholds	144
Log content	145
Limitations	146
Understand and optimize costs	147
What contributes to my costs?	147
What is the best way to lower my costs? How do I lower ingestion costs?	147
What is the best way to lower my query costs?	147
If I decrease the retention period of my metrics, will that help reduce my total bill?	148
How can I keep my alert query costs low?	148
What metrics can I use to monitor my costs?	149
Can I check my bill at any time?	149
Why is my bill higher at the beginning of the month than at the end of the month?	149
I deleted all my Amazon Managed Service for Prometheus workspaces, but I still seem to be	
getting charged. What might be happening?	150
Integrations	151
Amazon EKS cost monitoring	151
AWS Observability Accelerator	152
Prerequisites	152
Using the infrastructure monitoring example	153
AWS Controllers for Kubernetes	154
Prerequisites	155
Deploying a workspace	156
Configure cluster for remote write	160

Amazon CloudWatch metrics with Firehose	162
Infrastructure	162
Creating a Amazon CloudWatch stream	164
Cleanup	165
Security	167
Data protection	168
Data collected by Amazon Managed Service for Prometheus	169
Encryption at rest	169
Identity and Access Management	182
Audience	183
Authenticating with identities	184
Managing access using policies	187
How Amazon Managed Service for Prometheus works with IAM	189
Identity-based policy examples	196
Troubleshooting	199
IAM permissions and policies	201
Amazon Managed Service for Prometheus permissions	201
Sample IAM policies	201
Compliance Validation	202
Resilience	203
Infrastructure Security	203
Using service-linked roles	204
Metric scraping role	204
CloudTrail logs	206
Amazon Managed Service for Prometheus management events in CloudTrail	208
Amazon Managed Service for Prometheus event examples	208
Set up IAM roles for service accounts	213
Set up service roles for the ingestion of metrics from Amazon EKS clusters	213
Set up IAM roles for service accounts for the querying of metrics	216
Interface VPC endpoints	219
Create an interface VPC endpoint for Amazon Managed Service for Prometheus	220
Troubleshooting	223
429 or limit exceeded errors	223
I see duplicate samples	224
I see errors about sample timestamps	225
I see an error message related to a limit	225

Your local Prometheus server output exceeds the limit	226
Some of my data isn't appearing	227
Tagging	228
Tagging workspaces	229
Add a tag to a workspace	229
View tags for a workspace	231
Edit tags for a workspace	232
Remove a tag from a workspace	233
Tagging rule groups namespaces	235
Add a tag to a rule groups namespace	. 235
View tags for a rule groups namespace	237
Edit tags for a rule groups namespace	238
Remove a tag from a rule groups namespace	. 239
Service quotas	. 241
Service quotas	241
Active series default	247
Ingestion throttling	248
Additional limits on ingested data	249
API Reference	250
Amazon Managed Service for Prometheus APIs	250
Using Amazon Managed Service for Prometheus with an AWS SDK	250
Prometheus-compatible APIs	251
CreateAlertManagerAlerts	252
DeleteAlertManagerSilence	. 253
GetAlertManagerStatus	254
GetAlertManagerSilence	255
GetLabels	256
GetMetricMetadata	259
GetSeries	260
ListAlerts	262
ListAlertManagerAlerts	263
ListAlertManagerAlertGroups	264
ListAlertManagerReceivers	266
ListAlertManagerSilences	267
ListRules	269
PutAlertManagerSilences	270

QueryMetric	s
RemoteWrite	274
Document History	/

What is Amazon Managed Service for Prometheus?

Amazon Managed Service for Prometheus is a serverless, Prometheus-compatible monitoring service for container metrics that makes it easier to securely monitor container environments at scale. With Amazon Managed Service for Prometheus, you can use the same open-source Prometheus data model and query language that you use today to monitor the performance of your containerized workloads, and also enjoy improved scalability, availability, and security without having to manage the underlying infrastructure.

Amazon Managed Service for Prometheus automatically scales the ingestion, storage, and querying of operational metrics as workloads scale up and down. It integrates with AWS security services to enable fast and secure access to data.

Amazon Managed Service for Prometheus is designed to be highly available using multiple Availability Zone (Multi-AZ) deployments. Data ingested into a workspace is replicated across three Availability Zones in the same Region.

Amazon Managed Service for Prometheus works with container clusters that run on Amazon Elastic Kubernetes Service and self-managed Kubernetes environments.

With Amazon Managed Service for Prometheus, you use the same open-source Prometheus data model and PromQL query language that you use with Prometheus. Engineering teams can use PromQL to filter, aggregate, and alarm on metrics and quickly gain performance visibility without any code changes. Amazon Managed Service for Prometheus provides flexible query capabilities without the operational cost and complexity.

Metrics ingested into a workspace are stored for 150 days by default, and are then automatically deleted. You can adjust the retention period by configuring your workspace up to a maximum of 1095 days (three years). For more information, see <u>Configure your workspace</u>.

Supported Regions

Amazon Managed Service for Prometheus currently supports the following Regions:

Region Name	Region	Endpoint	Protocol
US East	us-east-2	aps.us-east-2.amazonaws.com	HTTPS
(Ohio)		aps-workspaces.us-east-2.amazonaws.com	HTTPS
		aps-workspaces-fips.us-east-2.amazon	HTTPS
		aws.com	HTTPS
		aps-workspaces-fips.us-east-2.api.aws	HTTPS
		aps-workspaces.us-east-2.api.aws	HTTPS
		aps-fips.us-east-2.amazonaws.com	HTTPS
		aps.us-east-2.api.aws	HTTPS
		aps-fips.us-east-2.api.aws	
US East (N	us-east-1	aps.us-east-1.amazonaws.com	HTTPS
East (N. Virginia)		aps-workspaces.us-east-1.amazonaws.com	HTTPS
		aps-workspaces-fips.us-east-1.amazon aws.com	HTTPS
			HTTPS
		aps-workspaces-fips.us-east-1.api.aws	HTTPS
		aps-workspaces.us-east-1.api.aws	HTTPS
		aps-fips.us-east-1.amazonaws.com	HTTPS
		aps.us-east-1.api.aws	HTTPS
		aps-fips.us-east-1.api.aws	
US West	US- wort 2	aps.us-west-2.amazonaws.com	HTTPS
(Oregon)	west-2	aps-workspaces.us-west-2.amazonaws.com	HTTPS
			HTTPS

Region	Endpoint	Protocol
	aps-workspaces-fips.us-west-2.amazon aws.com	HTTPS HTTPS
	aps-workspaces-fips.us-west-2.api.aws	HTTPS
	aps-fips.us-west-2.amazonaws.com	HTTPS
	aps.us-west-2.api.aws	HTTPS
	aps-fips.us-west-2.api.aws	
af-south- 1	aps.af-south-1.amazonaws.com	HTTPS
	aps-workspaces.af-south-1.amazonaws.com	HTTPS
		HTTPS HTTPS
ap- east-1	aps.ap-east-1.amazonaws.com	HTTPS
	aps-workspaces.ap-east-1.amazonaws.com	HTTPS
	aps-workspaces.ap-east-1.api.aws	HTTPS
	aps.ap-east-1.api.aws	HTTPS
ap- southe ast-5	aps.ap-southeast-5.amazonaws.com	HTTPS
	aps-workspaces.ap-southeast-5.amazon aws.com	HTTPS
	aps-workspaces.ap-southeast-5.api.aws	HTTPS
	aps.ap-southeast-5.api.aws	HTTPS
	af-south- 1 ap- east-1 ap- southe	Image: source series and ser

Region Name	Region	Endpoint	Protocol	
Asia Pacific (Mumbai)	ap- south-1	aps.ap-south-1.amazonaws.com	HTTPS	
		aps-workspaces.ap-south-1.amazonaws.com	HTTPS	
		aps-workspaces.ap-south-1.api.aws	HTTPS	
		aps.ap-south-1.api.aws	HTTPS	
Asia	ap-	aps.ap-northeast-2.amazonaws.com	HTTPS	
Pacific (Seoul)	northe ast-2	aps-workspaces.ap-northeast-2.amazon	HTTPS	
		aws.com	HTTPS	
		aps-workspaces.ap-northeast-2.api.aws	HTTPS	
		aps.ap-northeast-2.api.aws		
Asia Pacific (Singapor e)	ap- southe ast-1	aps.ap-southeast-1.amazonaws.com	HTTPS	
		ans-workspaces an-southeast-1 amazon	HTTPS	
			aps-workspaces.ap-southeast-1.api.aws	HTTPS
			HTTPS	
		aps.ap-southeast-1.api.aws		
Asia Docific	ap- southe ast-2	aps.ap-southeast-2.amazonaws.com	HTTPS	
Pacific (Sydney)		aps-workspaces.ap-southeast-2.amazon	HTTPS	
			aws.com	HTTPS
		aps-workspaces.ap-southeast-2.api.aws	HTTPS	
		aps.ap-southeast-2.api.aws		

Region Name	Region	Endpoint	Protocol	
Asia	ap-	aps.ap-southeast-7.amazonaws.com	HTTPS	
Pacific (Thailand	southe ast-7	aps-workspaces.ap-southeast-7.amazon	HTTPS	
)		aws.com	HTTPS	
		aps-workspaces.ap-southeast-7.api.aws aps.ap-southeast-7.api.aws	HTTPS	
Asia	ap-	aps.ap-northeast-1.amazonaws.com	HTTPS	
Pacific (Tokyo)	northe ast-1	aps-workspaces.ap-northeast-1.amazon	HTTPS	
		aws.com	HTTPS	
		aps-workspaces.ap-northeast-1.api.aws aps.ap-northeast-1.api.aws	HTTPS	
Canada	ca-centra	aps.ca-central-1.amazonaws.com	HTTPS	
(Central)	l-1	aps-workspaces.ca-central-1.amazonaws.com	HTTPS	
		aps-workspaces-fips.ca-central-1.ama zonaws.com aps-workspaces-fips.ca-central-1.api.aws		HTTPS
				HTTPS
				HTTPS
		aps-workspaces.ca-central-1.api.aws	HTTPS	
		aps-fips.ca-central-1.amazonaws.com	HTTPS	
		aps.ca-central-1.api.aws	HTTPS	
		aps-fips.ca-central-1.api.aws		

Region Name	Region	Endpoint	Protocol
Europe	eu-	aps.eu-central-1.amazonaws.com	HTTPS
(Frankfur t)	central-1	aps-workspaces.eu-central-1.amazonaws.com	HTTPS
		aps-workspaces.eu-central-1.api.aws	HTTPS
		aps.eu-central-1.api.aws	HTTPS
Europe	eu-	aps.eu-west-1.amazonaws.com	HTTPS
(Ireland)	west-1	aps-workspaces.eu-west-1.amazonaws.com	HTTPS
		aps-workspaces.eu-west-1.api.aws	HTTPS
		aps.eu-west-1.api.aws	HTTPS
Europe (London)	eu- west-2	aps.eu-west-2.amazonaws.com	HTTPS
(London)		aps-workspaces.eu-west-2.amazonaws.com	HTTPS
		aps-workspaces.eu-west-2.api.aws	HTTPS
		aps.eu-west-2.api.aws	HTTPS
Europe (Milan)	eu- south-1	aps.eu-south-1.amazonaws.com	HTTPS
(Pintan)		aps-workspaces.eu-south-1.amazonaws.com	HTTPS
		aps-workspaces.eu-south-1.api.aws	HTTPS
		aps.eu-south-1.api.aws	HTTPS
Europe (Paris)	eu- west-3	aps.eu-west-3.amazonaws.com	HTTPS
(1 4115)	inest 5	aps-workspaces.eu-west-3.amazonaws.com	HTTPS
		aps-workspaces.eu-west-3.api.aws	HTTPS
		aps.eu-west-3.api.aws	HTTPS

Region Name	Region	Endpoint	Protocol	
Europe (Stockhol m)	eu-	aps.eu-north-1.amazonaws.com	HTTPS	
	north-1	aps-workspaces.eu-north-1.amazonaws.com	HTTPS	
		aps-workspaces.eu-north-1.api.aws	HTTPS	
		aps.eu-north-1.api.aws	HTTPS	
Europe	eu-	aps.eu-central-2.amazonaws.com	HTTPS	
(Zurich)	central-2	aps-workspaces.eu-central-2.amazonaws.com	HTTPS	
		aps-workspaces.eu-central-2.api.aws	HTTPS	
		aps.eu-central-2.api.aws	HTTPS	
Middle East (UAE)	me- central-1	aps.me-central-1.amazonaws.com	HTTPS	
		aps-workspaces.me-central-1.amazonaws.com	HTTPS	
		aps-workspaces.me-central-1.api.aws	HTTPS	
		aps.me-central-1.api.aws	HTTPS	
South America	sa-east-1	aps.sa-east-1.amazonaws.com	HTTPS	
(São		aps-workspaces.sa-east-1.amazonaws.com	HTTPS	
Paulo)		aps-workspaces.sa-east-1.api.aws	HTTPS	
		aps.sa-east-1.api.aws	HTTPS	

Amazon Managed Service for Prometheus includes control plane endpoints (to perform workspace management tasks) and data plane endpoints (to work with Prometheus-compatible data in a workspace instance). Control plane endpoints start with aps.*, and dataplane endpoints start with aps-workspaces.*. Endpoints that end in .amazonaws.com support IPv4, and endpoints that end in .api.aws support both IPv4 and IPv6.

Pricing

You incur charges for ingestion and storage of metrics. Storage charges are based on the compressed size of metric samples and metadata. For more information, see <u>Amazon Managed</u> <u>Service for Prometheus Pricing</u>.

You can use AWS Cost Explorer and AWS Cost and Usage Reports to monitor your charges. For more information, see <u>Exploring your data using Cost Explorer</u> and <u>What are AWS Cost and Usage</u> <u>Reports</u>.

Premium support

If you subscribe to any level of the AWS premium support plans, your premium support applies to Amazon Managed Service for Prometheus.

Get started with Amazon Managed Service for Prometheus

Amazon Managed Service for Prometheus is a serverless, Prometheus-compatible service for monitoring container metrics that makes it easy to securely monitor container environments at scale. This section takes you through three key areas of using Amazon Managed Service for Prometheus:

- <u>Create a workspace</u> Create a Amazon Managed Service for Prometheus workspace to store and monitor your metrics.
- <u>Ingest metrics</u> Your workspace is empty until you get metrics into your workspace. You can send metrics to Amazon Managed Service for Prometheus, or have Amazon Managed Service for Prometheus scrape metrics automatically.
- <u>Query metrics</u> Once you have metrics as data in your workspace, you are ready to query the data to explore or monitor those metrics.

If you are new to AWS, this section also includes details about setting up an AWS account.

Topics

- Set up AWS
- Create an Amazon Managed Service for Prometheus workspace
- Ingest Prometheus metrics to the workspace
- Query your Prometheus metrics

Set up AWS

Complete the tasks in this section to get set up with AWS for the first time. If you already have an AWS account, skip ahead to <u>Create an Amazon Managed Service for Prometheus workspace</u>.

When you sign up for AWS, your AWS account automatically has access to all services in AWS, including Amazon Managed Service for Prometheus. However, you are charged only for the services that you use.

Topics

• Sign up for an AWS account

Sign up for an AWS account

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

- 1. Open https://portal.aws.amazon.com/billing/signup.
- 2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call or text message and entering a verification code on the phone keypad.

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform <u>tasks that require root</u> <u>user access</u>.

AWS sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <u>https://aws.amazon.com/</u> and choosing **My Account**.

Create a user with administrative access

After you sign up for an AWS account, secure your AWS account root user, enable AWS IAM Identity Center, and create an administrative user so that you don't use the root user for everyday tasks.

Secure your AWS account root user

1. Sign in to the <u>AWS Management Console</u> as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

For help signing in by using root user, see <u>Signing in as the root user</u> in the AWS Sign-In User Guide.

2. Turn on multi-factor authentication (MFA) for your root user.

For instructions, see Enable a virtual MFA device for your AWS account root user (console) in the IAM User Guide.

Create a user with administrative access

1. Enable IAM Identity Center.

For instructions, see Enabling AWS IAM Identity Center in the AWS IAM Identity Center User *Guide*.

2. In IAM Identity Center, grant administrative access to a user.

For a tutorial about using the IAM Identity Center directory as your identity source, see <u>Configure user access with the default IAM Identity Center directory</u> in the AWS IAM Identity Center User Guide.

Sign in as the user with administrative access

• To sign in with your IAM Identity Center user, use the sign-in URL that was sent to your email address when you created the IAM Identity Center user.

For help signing in using an IAM Identity Center user, see <u>Signing in to the AWS access portal</u> in the AWS Sign-In User Guide.

Assign access to additional users

1. In IAM Identity Center, create a permission set that follows the best practice of applying leastprivilege permissions.

For instructions, see Create a permission set in the AWS IAM Identity Center User Guide.

2. Assign users to a group, and then assign single sign-on access to the group.

For instructions, see <u>Add groups</u> in the AWS IAM Identity Center User Guide.

Create an Amazon Managed Service for Prometheus workspace

A *workspace* is a logical space dedicated to the storage and querying of Prometheus metrics. A workspace supports fine-grained access control for authorizing its management such as update, list, describe, and delete, and the ingestion and querying of metrics. You can have one or more workspaces in each Region in your account.

To set up a workspace, follow these steps.

í) Note

For more detailed information about creating a workspace and the options available, see Create a Amazon Managed Service for Prometheus workspace.

To create a Amazon Managed Service for Prometheus workspace

- 1. Open the Amazon Managed Service for Prometheus console at https://console.aws.amazon.com/prometheus/.
- 2. For **Workspace alias**, enter an alias for the new workspace.

Workspace aliases are friendly names that help you identify your workspaces. They do not have to be unique. Two workspaces could have the same alias, but all workspaces will have unique workspace IDs, which are generated by Amazon Managed Service for Prometheus.

3. (Optional) To add tags to the namespace, choose Add new tag.

Then, for Key, enter a name for the tag. You can add an optional value for the tag in Value.

To add another tag, choose **Add new tag** again.

4. Choose **Create workspace**.

The workspace details page appears. This displays information including the status, ARN, workspace ID, and endpoint URLs for this workspace for both remote write and queries.

Initially, the status is probably **CREATING**. Wait until the status is **ACTIVE before you move on** to setting up your metric ingestion.

Make notes of the URLs displayed for **Endpoint - remote write URL** and **Endpoint - query URL**. You'll need them when you configure your Prometheus server to remote write metrics to this workspace and when you query those metrics.

Ingest Prometheus metrics to the workspace

One way to ingest metrics is to use a standalone Prometheus *agent* (a Prometheus instance running in agent mode) to scrape metrics from your cluster and forward them to Amazon Managed Service for Prometheus for storage and monitoring. This section explains how to set up the

ingestion of metrics into your Amazon Managed Service for Prometheus workspace from Amazon EKS by setting up a new instance of Prometheus agent using Helm.

To generate metrics in Amazon EKS, such as Kubernetes or node-level metrics, you can use the Amazon EKS community add-ons. For more information, see <u>Available community add-ons</u> in the *Amazon EKS User Guide*.

For information about other ways to ingest data into Amazon Managed Service for Prometheus, including how to secure metrics and create high-availability metrics, see <u>Ingest metrics to your</u> Amazon Managed Service for Prometheus workspace.

🚯 Note

Metrics ingested into a workspace are stored for 150 days by default, and are then automatically deleted. You can adjust the retention period by configuring your workspace up to a maximum of 1095 days (three years). For more information, see <u>Configure your</u> <u>workspace</u>.

The instructions in this section get you up and running with Amazon Managed Service for Prometheus quickly. It assumes that you have already <u>created a workspace</u>. In this section, you set up a new Prometheus server in an Amazon EKS cluster, and the new server uses a default configuration to act as an agent to send metrics to Amazon Managed Service for Prometheus. This method has the following prerequisites:

- You must have an Amazon EKS cluster from which the new Prometheus server will collect metrics.
- Your Amazon EKS cluster must have an Amazon EBS CSI driver installed (required by Helm).
- You must use Helm CLI 3.0 or later.
- You must use a Linux or MacOS computer to perform the steps in the following sections.

Step 1: Add new Helm chart repositories

To add new Helm chart repositories, enter the following commands. For more information about these commands, see Helm Repo.

helm repo add prometheus-community https://prometheus-community.github.io/helm-charts

```
helm repo add kube-state-metrics https://kubernetes.github.io/kube-state-metrics
helm repo update
```

Step 2: Create a Prometheus namespace

Enter the following command to create a Prometheus namespace for the Prometheus server and other monitoring components. Replace *prometheus-agent-namespace* with the name that you want for this namespace.

```
kubectl create namespace prometheus-agent-namespace
```

Step 3: Set up IAM roles for service accounts

For this method of ingestion, you need to use IAM roles for service accounts in the Amazon EKS cluster where the Prometheus agent is running.

With IAM roles for service accounts, you can associate an IAM role with a Kubernetes service account. This service account can then provide AWS permissions to the containers in any pod that uses that service account. For more information, see IAM roles for service accounts.

If you have not already set up these roles, follow the instructions at <u>Set up service roles for</u> <u>the ingestion of metrics from Amazon EKS clusters</u> to set up the roles. The instructions in that section require the use of eksctl. For more information, see <u>Getting started with Amazon Elastic</u> <u>Kubernetes Service – eksctl</u>.

Note

When you are not on EKS or AWS and using just access key and secret key to access Amazon Managed Service for Prometheus, you cannot use the EKS-IAM-ROLE based SigV4.

Step 4: Set up the new server and start ingesting metrics

To install the new Prometheus agent and send metrics to your Amazon Managed Service for Prometheus workspace, follow these steps.

To install a new Prometheus agent and send metrics to your Amazon Managed Service for Prometheus workspace

- Use a text editor to create a file named my_prometheus_values_yaml with the following content.
 - Replace IAM_PROXY_PROMETHEUS_ROLE_ARN with the ARN of the amp-iamproxy-ingestrole that you created in <u>Set up service roles for the ingestion of metrics from Amazon EKS</u> clusters.
 - Replace WORKSPACE_ID with the ID of your Amazon Managed Service for Prometheus workspace.
 - Replace *REGION* with the Region of your Amazon Managed Service for Prometheus workspace.

```
## The following is a set of default values for prometheus server helm chart which
enable remoteWrite to AMP
## For the rest of prometheus helm chart values see: https://github.com/prometheus-
community/helm-charts/blob/main/charts/prometheus/values.yaml
##
serviceAccounts:
  server:
    name: amp-iamproxy-ingest-service-account
   annotations:
      eks.amazonaws.com/role-arn: ${IAM_PROXY_PROMETHEUS_ROLE_ARN}
server:
 remoteWrite:
    - url: https://aps-workspaces.${REGION}.amazonaws.com/workspaces/
${WORKSPACE_ID}/api/v1/remote_write
      sigv4:
        region: ${REGION}
      queue_config:
        max_samples_per_send: 1000
        max_shards: 200
        capacity: 2500
```

- 2. Enter the following command to create the Prometheus server.
 - Replace *prometheus-chart-name* with your Prometheus release name.
 - Replace *prometheus-agent-namespace* with the name of your Prometheus namespace.

```
helm install prometheus-chart-name prometheus-community/prometheus -n prometheus-
agent-namespace \
-f my_prometheus_values_yaml
```

Query your Prometheus metrics

Now that metrics are being ingested to the workspace, you can query them. A common way to query your metrics is to use a service such as Grafana to query the metrics. In this section, you will learn how to use Amazon Managed Grafana to query metrics from Amazon Managed Service for Prometheus.

1 Note

To learn about other ways to query your Amazon Managed Service for Prometheus metrics, or use the Amazon Managed Service for Prometheus APIs, see <u>Query your Prometheus</u> <u>metrics</u>.

This section assumes you already have a workspace created, and are ingesting metrics into it.

You perform your queries using the standard Prometheus query language, PromQL. For more information about PromQL and its syntax, see <u>Querying Prometheus</u> in the Prometheus documentation.

Amazon Managed Grafana is a fully managed service for open-source Grafana that simplifies connecting to open-source, third-party ISV, and AWS services for visualizing and analyzing your data sources at scale.

Amazon Managed Service for Prometheus supports using Amazon Managed Grafana to query metrics in a workspace. In the Amazon Managed Grafana console, you can add an Amazon Managed Service for Prometheus workspace as a data source by discovering your existing Amazon Managed Service for Prometheus accounts. Amazon Managed Grafana manages the configuration of the authentication credentials that are required to access Amazon Managed Service for Prometheus. For detailed instructions on creating a connection to Amazon Managed Service for Prometheus from Amazon Managed Grafana, see the instructions in <u>the Amazon Managed Grafana</u> User Guide. You may also view your Amazon Managed Service for Prometheus alerts in Amazon Managed Grafana. For instructions to set up integration with alerts, see <u>Integrate alerts with Amazon</u> Managed Grafana or open source Grafana.

(i) Note

If you have configured your Amazon Managed Grafana workspace to use a Private VPC, you must connect your Amazon Managed Service for Prometheus workspace to the same VPC. For more information, see <u>Connecting to Amazon Managed Grafana in a private VPC</u>.

Manage Amazon Managed Service for Prometheus workspaces

A *workspace* is a logical space dedicated to the storage and querying of Prometheus metrics. A workspace supports fine-grained access control for authorizing its management such as update, list, describe, and delete, and the ingestion and querying of metrics. You can have one or more workspaces in each Region in your account.

Use the procedures in this section to create and manage your Amazon Managed Service for Prometheus workspaces.

Topics

- Create a Amazon Managed Service for Prometheus workspace
- Configure your workspace
- Edit a workspace alias
- Find your Amazon Managed Service for Prometheus workspace details, including ARN
- Delete an Amazon Managed Service for Prometheus workspace

Create a Amazon Managed Service for Prometheus workspace

Follow these steps to create a Amazon Managed Service for Prometheus workspace. You can choose to use the AWS CLI or the Amazon Managed Service for Prometheus console.

🚯 Note

If you are running an Amazon EKS cluster, you can also create a new workspace using <u>AWS</u> <u>Controllers for Kubernetes</u>.

To create a workspace using the AWS CLI

 Enter the following command to create the workspace. This example creates a workspace named my-first-workspace, but you can use a different alias (or none) if you want. Workspace aliases are friendly names that help you identify your workspaces. They do not have to be unique. Two workspaces can have the same alias, but all workspaces have unique workspace IDs, which are generated by Amazon Managed Service for Prometheus. (Optional) To use your own KMS key to encrypt data stored in your workspace, you can include the kmsKeyArn parameter with the AWS KMS key to use. While Amazon Managed Service for Prometheus does not charge you for using customer managed keys, there may be costs associated with keys from AWS Key Management Service. For more information about Amazon Managed Service for Prometheus encryption of data in the workspace, or how to create, manage, and use your own customer managed key, see Encryption at rest.

Parameters in brackets ([]) are optional, do not include the brackets in your command.

```
aws amp create-workspace [--alias my-first-workspace] [--kmsKeyArn arn:aws:aps:us-
west-2:111122223333:workspace/ws-sample-1234-abcd-56ef-7890abcd12ef] [--
tags Status=Secret,Team=My-Team]
```

This command returns the following data:

- workspaceId is the unique ID for this workspace. Make a note of this ID.
- arn is the ARN for this workspace.
- status is the current status of the workspace. Immediately after you create the workspace, this will probably be CREATING.
- kmsKeyArn is the customer managed key used to encrypt the workspace data, if given.

🚺 Note

Workspaces created with customer managed keys cannot use <u>AWS managed</u> collectors for ingestion.

Choose whether to use customer managed keys or AWS owned keys carefully. Workspaces created with customer managed keys can't be converted to use AWS owned keys later (and vice versa).

- tags lists the workspace's tags, if any.
- 2. If your create-workspace command returns a status of CREATING, you can then enter the following command to determine when the workspace is ready. Replace *my-workspace-id* with the value that the create-workspace command returned for workspaceId.

```
aws amp describe-workspace --workspace-id my-workspace-id
```

When the describe-workspace command returns ACTIVE for status, the workspace is ready to use.

To create a workspace using the Amazon Managed Service for Prometheus console

- 1. Open the Amazon Managed Service for Prometheus console at <u>https://</u> console.aws.amazon.com/prometheus/.
- 2. Choose Create.
- 3. For **Workspace alias**, enter an alias for the new workspace.

Workspace aliases are friendly names that help you identify your workspaces. They do not have to be unique. Two workspaces can have the same alias, but all workspaces have unique workspace IDs, which are generated by Amazon Managed Service for Prometheus.

4. (Optional) To use your own KMS key to encrypt data stored in your workspace, you can select Customize encryption settings, and choose the AWS KMS key to use (or create a new one). You can choose a key in your account from the drop down list, or enter the ARN for any key that you have access to. While Amazon Managed Service for Prometheus does not charge you for using customer managed keys, there may be costs associated with keys from AWS Key Management Service.

For more information about Amazon Managed Service for Prometheus encryption of data in the workspace, or how to create, manage, and use your own, customer managed key, see Encryption at rest.

Note

Workspaces created with customer managed keys cannot use <u>AWS managed collectors</u> for ingestion.

Choose whether to use customer managed keys or AWS owned keys carefully. Workspaces created with customer managed keys can't be converted to use AWS owned keys later (and vice versa).

5. (Optional) To add one or more tags to the workspace, choose **Add new tag**. Then, in **Key**, enter a name for the tag. You can add an optional value for the tag in **Value**.

To add another tag, choose Add new tag again.

6. Choose Create workspace.

The workspace details page appears. This displays information including the status, ARN, workspace ID, and endpoint URLs for this workspace for both remote write and queries.

The status returns **CREATING** until the workspace is ready. Wait until the status is **ACTIVE** before you move on to setting up your metric ingestion.

Make note of the URLs that are displayed for **Endpoint - remote write URL** and **Endpoint - query URL**. You'll need them when you configure your Prometheus server to remote write metrics to this workspace and when you query those metrics.

For information about how to ingest metrics into the workspace, see <u>Ingest Prometheus metrics to</u> <u>the workspace</u>.

Configure your workspace

You can configure your workspace for the following:

• Define *label sets* and define limits on the active time series that match your defined label sets. A label set is a set of one or more *labels*, which are name/value pairs that help give context to time series metrics.

By defining label sets and setting active time series limits, you can limit spikes in one tenant or source to affect only that tenant or source. For example, if you set a 1,000,000 active time series limit on the label set team=A env=prod, then if the number of ingested time series that match that label set exceed the limit, then only the time series that match the label set are throttled. This way, other tenants or metric sources are unaffected.

For more information about labels in Prometheus, see Data Model.

• Set a retention period to define the number of days for the data to be retained in the workspace.

To configure your workspace

- 1. Open the Amazon Managed Service for Prometheus console at https://console.aws.amazon.com/prometheus/.
- 2. In the upper left corner of the page, choose the menu icon and then choose **All workspaces**.
- 3. Choose the **Workspace ID** of the workspace.

- 4. Choose the **Workspace configurations** tab.
- 5. To set the retention period for the workspace, choose **Edit** in the **Retention period** section. Then specify the new retention period in days. The maximum is 1095 days (three years).
- 6. To add or modify label sets and their active series limits, choose **Edit** in the **Label sets** section. Then do the following:
 - a. (Optional) Enter a value in **Default bucket limit** to set a limit on the maximum number of active time series that can be ingested in the workspace, counting only time series that don't match any defined label set.
 - b. To define a label set, enter an active time series limit for the new label set under **Active** series limit.

Then, enter a label and value for one label that will be used in the label set, and choose **Add label**.

- c. (Optional) To define another label set, choose **Add another label set** and repeat the previous steps.
- 7. When you are finished, choose **Save changes**.

Edit a workspace alias

You can edit a workspace to change its alias. To change the workspace alias using the AWS CLI, enter the following command.

aws amp update-workspace-alias --workspace-id *my-workspace-id* --alias "new-alias"

To edit a workspace using the Amazon Managed Service for Prometheus console

- 1. Open the Amazon Managed Service for Prometheus console at https://console.aws.amazon.com/prometheus/.
- 2. In the upper left corner of the page, choose the menu icon and then choose All workspaces.
- 3. Choose the workspace ID of the workspace that you want to edit, and then choose **Edit**.
- 4. Enter a new alias for the workspace and then choose **Save**.

Find your Amazon Managed Service for Prometheus workspace details, including ARN

You can find the details of your Amazon Managed Service for Prometheus workspace by using either the AWS console or the AWS CLI.

Console

To find your workspace details using the Amazon Managed Service for Prometheus console

- 1. Open the Amazon Managed Service for Prometheus console at https://console.aws.amazon.com/prometheus/.
- 2. In the upper left corner of the page, choose the menu icon and then choose **All workspaces**.
- 3. Choose the **Workspace ID** of the workspace. This will display details about your workspace, including:
 - Current status The status of your workspace, for example Active, is displayed under Status.
 - **ARN** The workspace ARN is displayed under **ARN**.
 - **ID** The workspace ID is displayed under **Workspace ID**.
 - URLs The console displays multiple URLs for the workspace, including the URLs for writing to or querying data from the workspace.

🚺 Note

By default, the URLs given are the IPv4 URLs. You can also use dualstack (IPv4 and IPv6 supported) URLs. These are the same, but are in the domain api.aws rather than the default amazonaws.com. For example, if you were to see the following (an IPv4 URL):

```
https://aps-workspaces.us-east-1.amazonaws.com/workspaces/ws-abcd1234-
ef56-7890-ab12-example/api/v1/remote_write
```

You could create a dualstack (including support for IPv6), URL as follows:

```
https://aps-workspaces.us-east-1.api.aws/workspaces/ws-abcd1234-
ef56-7890-ab12-example/api/v1/remote_write
```

Below this section are tabs with information about rules, alert manager, logs, configuration, and tags.

AWS CLI

To find your workspace details using the AWS CLI

The following command returns the details of the workspace. You must replace *myworkspace-id* with the workspace ID of the workspace for which you want the details.

aws amp describe-workspace --workspace-id my-workspace-id

This returns details about your workspace, including:

- **Current status** The status of your workspace, for example ACTIVE, is returned in the statusCode property.
- **ARN** The workspace ARN is returned in the arn property.
- URLs The AWS CLI returns the base URL for the workspace in the prometheusEndpoint property.

🚯 Note

By default, the URL returned is the IPv4 URL. You can also use a dualstack (IPv4 and IPv6 supported) URL in the domain api.aws rather than the default amazonaws.com. For example, if you were to see the following (an IPv4 URL):

```
https://aps-workspaces.us-east-1.amazonaws.com/workspaces/ws-abcd1234-
ef56-7890-ab12-example/
```

You could create a dualstack (including support for IPv6), URL as follows:

```
https://aps-workspaces.us-east-1.api.aws/workspaces/ws-abcd1234-ef56-7890-
ab12-example/
```

You can also create the remote write and query URLs for the workspace, by adding / $api/v1/remote_write or /api/v1/query, respectively.$

Delete an Amazon Managed Service for Prometheus workspace

Deleting a workspace deletes the data that has been ingested into it.

i Note

Deleting an Amazon Managed Service for Prometheus workspace does not automatically delete any AWS managed collectors that are scraping metrics and sending them to the workspace. For more information, see Find and delete scrapers.

To delete a workspace using the AWS CLI

Use the following command:

aws amp delete-workspace --workspace-id my-workspace-id

To delete a workspace using the Amazon Managed Service for Prometheus console

- 1. Open the Amazon Managed Service for Prometheus console at https://console.aws.amazon.com/prometheus/.
- 2. In the upper left corner of the page, choose the menu icon and then choose All workspaces.
- 3. Choose the workspace ID of the workspace that you want to delete, and then choose **Delete**.
- 4. Enter **delete** in the confirmation box, and choose **Delete**.

Ingest metrics to your Amazon Managed Service for Prometheus workspace

Metrics must be ingested into your Amazon Managed Service for Prometheus workspace before you can query or alert on those metrics. This section explains how to set up the ingestion of metrics into your workspace.

🚯 Note

Metrics ingested into a workspace are stored for 150 days by default, and are then automatically deleted. You can adjust the retention period by configuring your workspace up to a maximum of 1095 days (three years). For more information, see <u>Configure your</u> <u>workspace</u>.

There are two methods of ingesting metrics into your Amazon Managed Service for Prometheus workspace.

- Using an AWS managed collector Amazon Managed Service for Prometheus provides a fully-managed, agentless scraper to automatically *scrape* metrics from your Amazon Elastic Kubernetes Service (Amazon EKS) clusters. Scraping automatically pulls the metrics from Prometheus-compatible endpoints.
- Using a customer managed collector You have many options for managing your own collector. Two of the most common collectors to use are installing your own instance of Prometheus, running in agent mode, or using AWS Distro for OpenTelemetry. These are both described in detail in the following sections.

Collectors send metrics to Amazon Managed Service for Prometheus using Prometheus remote write functionality. You can directly send metrics to Amazon Managed Service for Prometheus by using Prometheus remote write in your own application. For more details about directly using remote write, and remote write configurations, see <u>remote_write</u> in the Prometheus documentation.

Topics

Ingest metrics with AWS managed collectors

Ingest metrics with AWS managed collectors

A common use case for Amazon Managed Service for Prometheus is to monitor Kubernetes clusters managed by Amazon Elastic Kubernetes Service (Amazon EKS). Kubernetes clusters, and many applications that run within Amazon EKS, automatically export their metrics for Prometheus-compatible scrapers to access.

1 Note

Amazon EKS exposes API server metrics, kube-controller-manager metrics, and kube-scheduler metrics in a cluster. Many other technologies and applications running in Kubernetes environments provide Prometheus-compatible metrics. For a list of well-documented exporters, see Exporters and integrations in the Prometheus documentation.

Amazon Managed Service for Prometheus provides a fully managed, agent less scraper, or *collector*, that automatically discovers and pulls Prometheus-compatible metrics. You don't have to manage, install, patch, or maintain agents or scrapers. An Amazon Managed Service for Prometheus collector provides reliable, stable, highly available, automatically scaled collection of metrics for your Amazon EKS cluster. Amazon Managed Service for Prometheus managed collectors work with Amazon EKS clusters, including EC2 and Fargate.

An Amazon Managed Service for Prometheus collector creates an Elastic Network Interface (ENI) per subnet specified when creating the scraper. The collector scrapes the metrics through these ENIs, and uses remote_write to push the data to your Amazon Managed Service for Prometheus workspace using a VPC endpoint. The scraped data never travels on the public internet.

The following topics provide more information about how to use an Amazon Managed Service for Prometheus collector in your Amazon EKS cluster, and about the collected metrics.

Topics

- Using an AWS managed collector
- What are Prometheus-compatible metrics?

Using an AWS managed collector

To use an Amazon Managed Service for Prometheus collector, you must create a scraper that discovers and pulls metrics in your Amazon EKS cluster.

- You can create a scraper as part of your Amazon EKS cluster creation. For more information about creating an Amazon EKS cluster, including creating a scraper, see <u>Creating an Amazon EKS</u> cluster in the *Amazon EKS User Guide*.
- You can create your own scraper, programmatically with the AWS API or by using the AWS CLI.

An Amazon Managed Service for Prometheus collector scrapes metrics that are Prometheuscompatible. For more information about Prometheus compatible metrics, see <u>What are</u> <u>Prometheus-compatible metrics</u>?. Amazon EKS clusters expose metrics for the API server. Amazon EKS clusters that are Kubernetes version 1.28 or above also expose metrics for the kubescheduler and kube-controller-manager. For more information, see <u>Fetch control plane raw</u> <u>metrics in Prometheus format</u> in the *Amazon EKS User Guide*.

🚯 Note

Scraping metrics from a cluster may incur charges for network usage. One way to optimize these costs is to configure your /metrics endpoint to compress the provided metrics (for example, with gzip), reducing the data that must be moved across the network. How to do this depends on the application or library providing the metrics. Some libraries gzip by default.

The following topics describe how to create, manage, and configure scrapers.

Topics

- Create a scraper
- Configuring your Amazon EKS cluster
- Find and delete scrapers
- Scraper configuration
- Troubleshooting scraper configuration
- Scraper limitations

An Amazon Managed Service for Prometheus collector consists of a scraper that discovers and collects metrics from an Amazon EKS cluster. Amazon Managed Service for Prometheus manages the scraper for you, giving you the scalability, security, and reliability that you need, without having to manage any instances, agents, or scrapers yourself.

There are three ways to create a scraper:

- A scraper is automatically created for you when you <u>create an Amazon EKS cluster through the</u> <u>Amazon EKS console</u> and choose to turn on Prometheus metrics.
- You can create a scraper from the Amazon EKS console for an existing cluster. Open the cluster in the Amazon EKS console, then, on the **Observability** tab, choose **Add scraper**.

For more details on the available settings, see <u>Turn on Prometheus metrics</u> in the *Amazon EKS User Guide*.

• You can create a scraper using either the AWS API or the AWS CLI.

These options are described in the following procedure.

There are a few prerequisites for creating your own scraper:

- You must have an Amazon EKS cluster created.
- Your Amazon EKS cluster must have <u>cluster endpoint access control</u> set to include private access.
 It can include private and public, but must include private.
- The Amazon VPC in which the Amazon EKS cluster resides must have DNS enabled.

1 Note

The cluster will be associated with the scraper by its Amazon resource name (ARN). If you delete a cluster, and then create a new one with the same name, the ARN will be reused for the new cluster. Because of this, the scraper will attempt to collect metrics for the new cluster. You <u>delete scrapers</u> separately from deleting the cluster.

AWS API

To create a scraper using the AWS API

Use the CreateScraper API operation to create a scraper with the AWS API. The following example creates a scraper in the us-west-2 Region. You need to replace the AWS account, workspace, security, and Amazon EKS cluster information with your own IDs, and provide the configuration to use for your scraper.

Note

The security group and subnets should be set to the security group and subnets for the cluster to which you are connecting.

You must include at least two subnets, in at least two availability zones.

The scrapeConfiguration is a Prometheus configuration YAML file that is base64 encoded. You can download a general purpose configuration with the GetDefaultScraperConfiguration API operation. For more information about the format of the scrapeConfiguration, see Scraper configuration.

```
POST /scrapers HTTP/1.1
Content-Length: 415
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: aws-cli/1.18.147 Python/2.7.18 Linux/5.4.58-37.125.amzn2int.x86_64
 botocore/1.18.6
{
    "alias": "myScraper",
    "destination": {
        "ampConfiguration": {
            "workspaceArn": "arn:aws:aps:us-west-2:account-id:workspace/
ws-workspace-id"
        }
    },
    "source": {
        "eksConfiguration": {
            "clusterArn": "arn:aws:eks:us-west-2:account-id:cluster/cluster-name",
            "securityGroupIds": ["sg-security-group-id"],
            "subnetIds": ["subnet-subnet-id-1", "subnet-subnet-id-2"]
        }
    },
    "scrapeConfiguration": {
        "configurationBlob": <base64-encoded-blob>
```

```
}
```

}

AWS CLI

To create a scraper using the AWS CLI

Use the create-scraper command to create a scraper with the the AWS CLI. The following example creates a scraper in the us-west-2 Region. You need to replace the AWS account, workspace, security, and Amazon EKS cluster information with your own IDs, and provide the configuration to use for your scraper.

🚯 Note

The security group and subnets should be set to the security group and subnets for the cluster to which you are connecting.

You must include at least two subnets, in at least two availability zones.

The scrape-configuration is a Prometheus configuration YAML file that is base64 encoded. You can download a general purpose configuration with the get-default-scraper-configuration command. For more information about the format of the scrape-configuration, see <u>Scraper configuration</u>.

```
aws amp create-scraper \
    --source eksConfiguration="{clusterArn='arn:aws:eks:us-west-2:account-
id:cluster/cluster-name', securityGroupIds=['sg-security-group-
id'],subnetIds=['subnet-subnet-id-1', 'subnet-subnet-id-2']}" \
    --scrape-configuration configurationBlob=<base64-encoded-blob> \
    --destination ampConfiguration="{workspaceArn='arn:aws:aps:us-west-2:account-
id:workspace/ws-workspace-id'}"
```

The following is a full list of the scraper operations that you can use with the AWS API:

- Create a scraper with the <u>CreateScraper</u> API operation.
- List your existing scrapers with the ListScrapers API operation.
- Update the alias, configuration, or destination of a scraper with the <u>UpdateScraper</u> API operation.
- Delete a scraper with the <u>DeleteScraper</u> API operation.

- Get more details about a scraper with the <u>DescribeScraper</u> API operation.
- Get a general purpose configuration for scrapers with the <u>GetDefaultScraperConfiguration</u> API operation.

1 Note

The Amazon EKS cluster that you are scraping must be configured to allow Amazon Managed Service for Prometheus to access the metrics. The next topic describes how to configure your cluster.

Cross-account setup

To create a scraper in a cross-account setup when your Amazon EKS cluster from which you want to collect metrics is in a different account from the Amazon Managed Service for Prometheus collector, use the procedure below.

For example, when you have two accounts, the first source account account_id_source where the Amazon EKS is located, and a second target account account_id_target where the Amazon Managed Service for Prometheus workspace resides.

To create a scraper in a cross-account setup

 In the source account, create a role arn:aws:iam::account_id_source:role/Source and add the following trust policy.

```
}
}
```

2. On every combination of source (Amazon EKS cluster) and target (Amazon Managed Service for Prometheus workspace), you need to create a role arn:aws:iam::account_id_target:role/Target and add the following trust policy with permissions for AmazonPrometheusRemoteWriteAccess.


```
{
   "Effect": "Allow",
   "Principal": {
      "AWS": "arn:aws:iam::account_id_source:role/Source"
   },
   "Action": "sts:AssumeRole",
   "Condition": {
      "StringEquals": {
        "sts:ExternalId": "scraper_ARN"
      }
  }
}
```

3. Create a scraper with the --role-configuration option.

```
aws amp create-scraper \
    --source eksConfiguration="{clusterArn='arn:aws:eks:us-west-2:account-
id_source:cluster/xarw,subnetIds=[subnet-subnet-id]}" \
    --scrape-configuration configurationBlob=<base64-encoded-blob> \
    --destination ampConfiguration="{workspaceArn='arn:aws:aps:us-west-2:account-
id_target:workspace/ws-workspace-id'}"\
    --role-configuration '{"sourceRoleArn":"arn:aws:iam::account-id_source:role/
Source", "targetRoleArn":"arn:aws:iam::account-id_target:role/Target"}'
```

4. Validate the scraper creation.

```
aws amp list-scrapers
{
    "scrapers": [
        {
```



Changing between RoleConfiguration and service-linked role

When you want to switch back to a service-linked role instead of the RoleConfiguration to write to an Amazon Managed Service for Prometheus workspace, you must update the UpdateScraper and provide a workspace in the same account as the scraper without the

RoleConfiguration. The RoleConfiguration will be removed from the scraper and the service-linked role will be used.

When you are changing workspaces in the same account as the scraper and you want to continue using the RoleConfiguration, you must again provide the RoleConfiguration on UpdateScraper.

Creating scraper for workspaces enabled with customer managed keys

To create a scraper for ingesting metrics into a Amazon Managed Service for Prometheus workspace with <u>customer managed keys</u>, use the --role-configuration with both the source and target set to the same account.

```
aws amp create-scraper \
    --source eksConfiguration="{clusterArn='arn:aws:eks:us-west-2:account-id:cluster/
xarw,subnetIds=[subnet-subnet_id]}" \
    --scrape-configuration configurationBlob=<base64-encoded-blob> \
    --destination ampConfiguration="{workspaceArn='arn:aws:aps:us-west-2:account-
id:workspace/ws-workspace-id'}"\
    --role-configuration '{"sourceRoleArn":"arn:aws:iam::account_id:role/Source",
    "targetRoleArn":"arn:aws:iam::account_id:role/Target"}'
```

Common errors when creating scrapers

The following are the most common issues when attempting to create a new scraper.

- Required AWS resources don't exist. The *security group*, *subnets*, and *Amazon EKS cluster* specified must exist.
- Insufficient IP address space. You must have at least one IP address available in each subnet that you pass into the CreateScraper API.

Configuring your Amazon EKS cluster

Your Amazon EKS cluster must be configured to allow the scraper to access metrics. There are two options for this configuration:

- Use Amazon EKS *access entries* to automatically provide Amazon Managed Service for Prometheus collectors access to your cluster.
- Manually configure your Amazon EKS cluster for managed metric scraping.

The following topics describe each of these in more detail.

Configure Amazon EKS for scraper access with access entries

Using access entries for Amazon EKS is the easiest way to give Amazon Managed Service for Prometheus access to scrape metrics from your cluster.

The Amazon EKS cluster that you are scraping must be configured to allow API authentication. The cluster authentication mode must be set to either API or API_AND_CONFIG_MAP. This is viewable in the Amazon EKS console on the **Access configuration** tab of the cluster details. For more information, see <u>Allowing IAM roles or users access to Kubernetes object on your Amazon</u> <u>EKS cluster</u> in the *Amazon EKS User Guide*.

You can create the scraper when creating the cluster, or after creating the cluster:

- When creating a cluster You can configure this access when you create an Amazon EKS cluster through the Amazon EKS console (follow the instructions to create a scraper as part of the cluster), and an access entry policy will automatically be created, giving Amazon Managed Service for Prometheus access to the cluster metrics.
- Adding after a cluster is created if your Amazon EKS cluster already exists, then set the authentication mode to either API or API_AND_CONFIG_MAP, and any scrapers you create through the Amazon Managed Service for Prometheus API or CLI or through the Amazon EKS console will automatically have the correct access entry policy created for you, and the scrapers will have access to your cluster.

Access entry policy created

When you create a scraper and let Amazon Managed Service for Prometheus generate an access entry policy for you, it generates the following policy. For more information about access entries, see <u>Allowing IAM roles or users access to Kubernetes</u> in the *Amazon EKS User Guide*.

```
"nodes/proxy",
        "nodes/metrics",
        "services",
        "endpoints",
        "pods",
        "ingresses",
        "configmaps"
    ],
    "verbs": [
        "get",
        "list",
        "watch"
    ]
},
{
    "effect": "allow",
    "apiGroups": [
        "extensions",
        "networking.k8s.io"
    ],
    "resources": [
        "ingresses/status",
        "ingresses"
    ],
    "verbs": [
        "get",
        "list",
        "watch"
    ]
},
{
    "effect": "allow",
    "apiGroups": [
        "metrics.eks.amazonaws.com"
    ],
    "resources": [
        "kcm/metrics",
        "ksh/metrics"
    ],
    "verbs": [
        "get"
    ]
},
```

{

Manually configuring Amazon EKS for scraper access

If you prefer to use the aws-auth ConfigMap to control access to your kubernetes cluster, you can still give Amazon Managed Service for Prometheus scrapers access to your metrics. The following steps will give Amazon Managed Service for Prometheus access to scrape metrics from your Amazon EKS cluster.

i Note

For more information about ConfigMap and access entries, see <u>Allowing IAM roles or users</u> access to <u>Kubernetes</u> in the *Amazon EKS User Guide*.

This procedure uses kubectl and the AWS CLI. For information about installing kubectl, see Installing kubectl in the Amazon EKS User Guide.

To manually configure your Amazon EKS cluster for managed metric scraping

1. Create a file, called clusterrole-binding.yml, with the following text:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
   name: aps-collector-role
rules:
    - apiGroups: [""]
    resources: ["nodes", "nodes/proxy", "nodes/metrics", "services", "endpoints",
   "pods", "ingresses", "configmaps"]
    verbs: ["describe", "get", "list", "watch"]
    - apiGroups: ["extensions", "networking.k8s.io"]
```

```
resources: ["ingresses/status", "ingresses"]
   verbs: ["describe", "get", "list", "watch"]
  - nonResourceURLs: ["/metrics"]
    verbs: ["get"]
  - apiGroups: ["metrics.eks.amazonaws.com"]
    resources: ["kcm/metrics", "ksh/metrics"]
    verbs: ["get"]
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: aps-collector-user-role-binding
subjects:
- kind: User
  name: aps-collector-user
  apiGroup: rbac.authorization.k8s.io
roleRef:
  kind: ClusterRole
  name: aps-collector-role
  apiGroup: rbac.authorization.k8s.io
```

2. Run the following command in your cluster:

kubectl apply -f clusterrole-binding.yml

This will create the cluster role binding and rule. This example uses aps-collector-role as the role name, and aps-collector-user as the user name.

3. The following command gives you information about the scraper with the ID *scraper-id*. This is the scraper that you created using the command in the previous section.

aws amp describe-scraper --scraper-id scraper-id

4. From the results of the describe-scraper, find the roleArn.This will have the following format:

```
arn:aws:iam::account-id:role/aws-service-role/scraper.aps.amazonaws.com/
AWSServiceRoleForAmazonPrometheusScraper_unique-id
```

Amazon EKS requires a different format for this ARN. You must adjust the format of the returned ARN to be used in the next step. Edit it to match this format:

arn:aws:iam::account-id:role/AWSServiceRoleForAmazonPrometheusScraper_unique-id

For example, this ARN:

```
arn:aws:iam::111122223333:role/aws-service-role/scraper.aps.amazonaws.com/
AWSServiceRoleForAmazonPrometheusScraper_1234abcd-56ef-7
```

Must be rewritten as:

```
arn:aws:iam::111122223333:role/
AWSServiceRoleForAmazonPrometheusScraper_1234abcd-56ef-7
```

5. Run the following command in your cluster, using the modified roleArn from the previous step, as well as your cluster name and region.:

```
eksctl create iamidentitymapping --cluster cluster-name --region region-id --
arn roleArn --username aps-collector-user
```

This allows the scraper to access the cluster using the role and user you created in the clusterrole-binding.yml file.

Find and delete scrapers

You can use the AWS API or the AWS CLI to list the scrapers in your account or to delete them.

Note

Make sure that you are using the latest version of the AWS CLI or SDK. The latest version provides you with the latest features and functionality, as well as security updates. Alternatively, use <u>AWS Cloudshell</u>, which provides an always up-to-date command line experience, automatically.

To list all the scrapers in your account, use the ListScrapers API operation.

Alternatively, with the AWS CLI, call:

aws amp list-scrapers

ListScrapers returns all of the scrapers in your account, for example:

```
{
    "scrapers": [
        {
            "scraperId": "s-1234abcd-56ef-7890-abcd-1234ef567890",
            "arn": "arn:aws:aps:us-west-2:123456789012:scraper/s-1234abcd-56ef-7890-
abcd-1234ef567890",
            "roleArn": "arn:aws:iam::123456789012:role/aws-service-role/
AWSServiceRoleForAmazonPrometheusScraper_1234abcd-2931",
            "status": {
                "statusCode": "DELETING"
            },
            "createdAt": "2023-10-12T15:22:19.014000-07:00",
            "lastModifiedAt": "2023-10-12T15:55:43.487000-07:00",
            "tags": {},
            "source": {
                "eksConfiguration": {
                    "clusterArn": "arn:aws:eks:us-west-2:123456789012:cluster/my-
cluster",
                    "securityGroupIds": [
                        "sg-1234abcd5678ef90"
                    ],
                    "subnetIds": [
                        "subnet-abcd1234ef567890",
                        "subnet-1234abcd5678ab90"
                    ]
                }
            },
            "destination": {
                "ampConfiguration": {
                    "workspaceArn": "arn:aws:aps:us-west-2:123456789012:workspace/
ws-1234abcd-5678-ef90-ab12-cdef3456a78"
                }
            }
        }
    ]
}
```

To delete a scraper, find the scraperId for the scraper that you want to delete, using the ListScrapers operation, and then use the <u>DeleteScraper</u> operation to delete it.

Alternatively, with the AWS CLI, call:

```
aws amp delete-scraper --scraper-id scraperId
```

Scraper configuration

You can control how your scraper discovers and collects metrics with a Prometheus-compatible scraper configuration. For example, you can change the interval that metrics are sent to the workspace. You can also use relabeling to dynamically rewrite the labels of a metric. The scraper configuration is a YAML file that is part of the definition of the scraper.

When a new scraper is created, you specify a configuration by providing a base64 encoded YAML file in the API call. You can download a general purpose configuration file with the GetDefaultScraperConfiguration operation in the Amazon Managed Service for Prometheus API.

To modify the configuration of a scraper, you can use the UpdateScraper operation. If you need to update the source of the metrics (for example, to a different Amazon EKS cluster), you must delete the scraper and recreate it with the new source.

Supported configuration

For information about the scraper configuration format, including a detailed breakdown of the possible values, see <u>Configuration</u> in the Prometheus documentation. The global configuration options, and <scrape_config> options describe the most commonly needed options.

Because Amazon EKS is the only supported service, the only service discovery config (<*_sd_config>) supported is the <kubernetes_sd_config>.

The complete list of config sections allowed:

- <global>
- <scrape_config>
- <static_config>
- <relabel_config>

- <metric_relabel_configs>
- <kubernetes_sd_config>

Limitations within these sections are listed after the sample configuration file.

Sample configuration file

The following is a sample YAML configuration file with a 30 second scrape interval. This sample includes support for the kube API server metrics, as well as kube-controller-manager and kube-scheduler metrics. For more information, see <u>Fetch control plane raw metrics in Prometheus format</u> in the *Amazon EKS User Guide*.

```
global:
   scrape_interval: 30s
   external_labels:
     clusterArn: apiserver-test-2
scrape_configs:
  - job_name: pod_exporter
    kubernetes_sd_configs:
      - role: pod
  - job_name: cadvisor
    scheme: https
    authorization:
      type: Bearer
      credentials_file: /var/run/secrets/kubernetes.io/serviceaccount/token
    kubernetes_sd_configs:
      - role: node
    relabel_configs:
      - action: labelmap
        regex: __meta_kubernetes_node_label_(.+)
      - replacement: kubernetes.default.svc:443
        target_label: __address___
      - source_labels: [__meta_kubernetes_node_name]
        regex: (.+)
        target_label: __metrics_path___
        replacement: /api/v1/nodes/$1/proxy/metrics/cadvisor
  # apiserver metrics
  - scheme: https
    authorization:
      type: Bearer
      credentials_file: /var/run/secrets/kubernetes.io/serviceaccount/token
    job_name: kubernetes-apiservers
```

kubernetes_sd_configs:

```
User Guide
```

```
- role: endpoints
  relabel_configs:
  - action: keep
    regex: default;kubernetes;https
    source_labels:
    - ___meta_kubernetes_namespace
    - ___meta_kubernetes_service_name
    - __meta_kubernetes_endpoint_port_name
# kube proxy metrics
- job_name: kube-proxy
  honor_labels: true
  kubernetes_sd_configs:
  - role: pod
  relabel_configs:
  - action: keep
    source_labels:
    - __meta_kubernetes_namespace
    - __meta_kubernetes_pod_name
    separator: '/'
    regex: 'kube-system/kube-proxy.+'
  - source_labels:
    - __address___
    action: replace
    target_label: __address__
    regex: (.+?)(\\:\\d+)?
    replacement: $1:10249
# Scheduler metrics
- job_name: 'ksh-metrics'
  kubernetes_sd_configs:
  - role: endpoints
  metrics_path: /apis/metrics.eks.amazonaws.com/v1/ksh/container/metrics
  scheme: https
  bearer_token_file: /var/run/secrets/kubernetes.io/serviceaccount/token
  relabel_configs:
  - source_labels:
    - __meta_kubernetes_namespace
    - __meta_kubernetes_service_name
    - __meta_kubernetes_endpoint_port_name
    action: keep
    regex: default;kubernetes;https
# Controller Manager metrics
- job_name: 'kcm-metrics'
```

kubernetes_sd_configs:

```
- role: endpoints
metrics_path: /apis/metrics.eks.amazonaws.com/v1/kcm/container/metrics
scheme: https
bearer_token_file: /var/run/secrets/kubernetes.io/serviceaccount/token
relabel_configs:
- source_labels:
- __meta_kubernetes_namespace
- __meta_kubernetes_service_name
- __meta_kubernetes_endpoint_port_name
action: keep
regex: default;kubernetes;https
```

The following are limitations specific to AWS managed collectors:

- Scrape interval The scraper config can't specify a scrape interval of less than 30 seconds.
- Targets Targets in the static_config must be specified as IP addresses.
- DNS resolution Related to the target name, the only server name that is recognized in this config is the Kubernetes api server, kubernetes.default.svc. All other machines names must be specified by IP address.
- Authorization Omit if no authorization is needed. If it is needed, the authorization must be Bearer, and must point to the file /var/run/secrets/kubernetes.io/serviceaccount/ token. In other words, if used, the authorization section must look like the following:

```
authorization:
  type: Bearer
  credentials_file: /var/run/secrets/kubernetes.io/serviceaccount/token
```

Note

type: Bearer is the default, so can be omitted.

Troubleshooting scraper configuration

Amazon Managed Service for Prometheus collectors automatically discover and scrape metrics. But how can you troubleshoot when you don't see a metric you expect to see in your Amazon Managed Service for Prometheus workspace?

🔥 Important

Verify that private access for your Amazon EKS cluster is enabled. For more information, see Cluster private endpoint in the Amazon EKS User Guide.

The up metric is a helpful tool. For each endpoint that an Amazon Managed Service for Prometheus collector discovers, it automatically vends this metric. There are three states of this metric that can help you to troubleshoot what is happening within the collector.

 up is not present – If there is no up metric present for an endpoint, then that means that the collector was not able to find the endpoint.

If you are sure that the endpoint exists, there are several reasons why the collector might not be able to find it.

- You might need to adjust the scrape configuration. The discovery relabel_config might need to be adjusted.
- There could be a problem with the role used for discovery.
- The Amazon VPC used by the Amazon EKS cluster might not have <u>DNS enabled</u>, which would keep the collector from finding the endpoint.
- up is present, but is always 0 If up is present, but 0, then the collector is able to discover the endpoint, but can't find any Prometheus-compatible metrics.

In this case, you might try using a curl command against the endpoint directly. You can validate that you have the details correct, for example, the protocol (http or https), the endpoint, or port that you are using. You can also check that the endpoint is responding with a valid 200 response, and follows the Prometheus format. Finally, the body of the response can't be larger than the maximum allowed size. (For limits on AWS managed collectors, see the following section.)

 up is present and greater than 0 – If up is present, and is greater than 0, then metrics are being sent to Amazon Managed Service for Prometheus.

Validate that you are looking for the correct metrics in Amazon Managed Service for Prometheus (or your alternate dashboard, such as Amazon Managed Grafana). You can use curl again to check for expected data in your /metrics endpoint. Also check that you haven't exceeded other limits, such as the number of endpoints per scraper. You can check the number of metrics endpoints being scraped by checking the count of up metrics, using count(up).

Scraper limitations

There are few limitations to the fully managed scrapers provided by Amazon Managed Service for Prometheus.

- **Region** Your EKS cluster, managed scraper, and Amazon Managed Service for Prometheus workspace must all be in the same AWS Region.
- Account Your EKS cluster, managed scraper, and Amazon Managed Service for Prometheus workspace must all be in the same AWS account.
- **Collectors** You can have a maximum of 10 Amazon Managed Service for Prometheus scrapers per region per account.

🚯 Note

You can request an increase to this limit by requesting a quota increase.

- Metrics response The body of a response from any one /metrics endpoint request cannot be more than 50 megabytes (MB).
- Endpoints per scraper A scraper can scrape a maximum of 30,000 /metrics endpoints.
- Scrape interval The scraper config can't specify a scrape interval of less than 30 seconds.

What are Prometheus-compatible metrics?

To scrape Prometheus metrics from your applications and infrastructure for use in Amazon Managed Service for Prometheus, they must instrument and expose *Prometheus-compatible metrics* from Prometheus-compatible /metrics endpoints. You can implement your own metrics, but you don't have to. Kubernetes (including Amazon EKS) and many other libraries and services implement these metrics directly.

When metrics in Amazon EKS are exported to a Prometheus-compatible endpoint, you can have those metrics automatically scraped by the Amazon Managed Service for Prometheus collector.

For more information, see the following topics:

- For more information about existing libraries and services that export metrics as Prometheus metrics, see <u>Exporters and integrations</u> in the Prometheus documentation.
- For more information about exporting Prometheus-compatible metrics from your own code, see <u>Writing exporters</u> in the Prometheus documentation.

 For more information about how to set up an Amazon Managed Service for Prometheus collector to scrape metrics from your Amazon EKS clusters automatically, see <u>Using an AWS managed</u> collector.

Customer managed collectors

This section contains information about ingesting data by setting up your own collectors that send metrics to Amazon Managed Service for Prometheus using Prometheus remote write.

When you use your own collectors to send metrics to Amazon Managed Service for Prometheus, you are responsible for securing your metrics and making sure that the ingestion process meets your availability needs.

Most customer managed collectors use one of the following tools:

- AWS Distro for OpenTelemetry (ADOT) ADOT is a fully supported, secure, production-ready open source distribution of OpenTelemetry that provides agents to collect metrics. You can use ADOT to collect metrics and send them to your Amazon Managed Service for Prometheus workspace. For more information about the ADOT Collector, see AWS Distro for OpenTelemetry.
- Prometheus agent You can set up your own instance of the open source Prometheus server, running as an agent, to collect metrics and forward them to your Amazon Managed Service for Prometheus workspace.

The following topics describe using both of these tools and include general information about setting up your own collectors.

Topics

- Secure the ingestion of your metrics
- Using AWS Distro for OpenTelemetry as a collector
- Using a Prometheus instance as a collector
- Set up Amazon Managed Service for Prometheus for high availability data

Secure the ingestion of your metrics

Amazon Managed Service for Prometheus provides ways of helping you secure the ingestion of your metrics.

Using AWS PrivateLink with Amazon Managed Service for Prometheus

The network traffic of ingesting the metrics into Amazon Managed Service for Prometheus can be done over a public internet endpoint, or by a VPC endpoint through AWS PrivateLink. Using AWS PrivateLink ensures that the network traffic from your VPCs is secured within the AWS network without going over the public internet. To create an AWS PrivateLink VPC endpoint for Amazon Managed Service for Prometheus, see <u>Using Amazon Managed Service for Prometheus with</u> interface VPC endpoints.

Authentication and authorization

AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources. You use IAM to control who is authenticated (signed in) and authorized (has permissions) to use resources. Amazon Managed Service for Prometheus integrates with IAM to help you keep your data secure. When you set up Amazon Managed Service for Prometheus, you need to create some IAM roles that enable it to ingest metrics from Prometheus servers, and that enable Grafana servers to query the metrics that are stored in your Amazon Managed Service for Prometheus workspaces. For more information about IAM, see <u>What is IAM?</u>.

Another AWS security feature that can help you set up Amazon Managed Service for Prometheus is the AWS Signature Version 4 signing process (AWS SigV4). Signature Version 4 is the process to add authentication information to AWS requests sent by HTTP. For security, most requests to AWS must be signed with an access key, which consists of an access key ID and secret access key. These two keys are commonly referred to as your security credentials. For more information about SigV4, see <u>Signature Version 4 signing process</u>.

Using AWS Distro for OpenTelemetry as a collector

This section describes how to configure the AWS Distro for OpenTelemetry (ADOT) Collector to scrape from a Prometheus-instrumented application, and send the metrics to Amazon Managed Service for Prometheus. For more information about the ADOT Collector, see <u>AWS Distro for</u> <u>OpenTelemetry</u>.

The following topics describe three different ways to set up ADOT as a collector for your metrics, based on whether your metrics are coming from Amazon EKS, Amazon ECS, or an Amazon EC2 instance.

Topics

- <u>Set up metrics ingestion using AWS Distro for OpenTelemetry on an Amazon Elastic Kubernetes</u> Service cluster
- Set up metrics ingestion from Amazon ECS using AWS Distro for Open Telemetry
- Set up metrics ingestion from an Amazon EC2 instance using remote write

Set up metrics ingestion using AWS Distro for OpenTelemetry on an Amazon Elastic Kubernetes Service cluster

You can use the AWS Distor for OpenTelemetry (ADOT) collector to scrape metrics from a Prometheus-instrumented application, and send the metrics to Amazon Managed Service for Prometheus.

🚯 Note

For more information about the ADOT collector, see <u>AWS Distro for OpenTelemetry</u>. For more information about Prometheus-instrumented applications, see <u>What are</u> <u>Prometheus-compatible metrics</u>?

Collecting Prometheus metrics with ADOT involves three OpenTelemetry components: the Prometheus Receiver, the Prometheus Remote Write Exporter, and the Sigv4 Authentication Extension.

You can configure the Prometheus Receiver using your existing Prometheus configuration to perform service discovery and metric scraping. The Prometheus Receiver scrapes metrics in the Prometheus exposition format. Any applications or endpoints that you want to scrape should be configured with the Prometheus client library. The Prometheus Receiver supports the full set of Prometheus scraping and re-labeling configurations described in <u>Configuration</u> in the Prometheus documentation. You can paste these configurations directly into your ADOT Collector configurations.

The Prometheus Remote Write Exporter uses the remote_write endpoint to send the scraped metrics to your management portal workspace. The HTTP requests to export data will be signed with AWS SigV4, the AWS protocol for secure authentication, with the Sigv4 Authentication Extension. For more information, see <u>Signature Version 4 signing process</u>.

The collector automatically discovers Prometheus metrics endpoints on Amazon EKS and uses the configuration found in <<u>kubernetes_sd_config></u>.

The following demo is an example of this configuration on a cluster running Amazon Elastic Kubernetes Service or self-managed Kubernetes. To perform these steps, you must have AWS credentials from any of the potential options in the default AWS credentials chain. For more information, see <u>Configuring the AWS SDK for Go</u>. This demo uses a sample app that is used for integration tests of the process. The sample app exposes metrics at the /metrics endpoint, like the Prometheus client library.

Prerequisites

Before you begin the following ingestion setup steps, you must set up your IAM role for the service account and trust policy.

To set up the IAM role for service account and trust policy

1. Create the IAM role for the service account by following the steps in <u>Set up service roles for</u> the ingestion of metrics from Amazon EKS clusters.

The ADOT Collector will use this role when it scrapes and exports metrics.

- 2. Next, edit the trust policy. Open the IAM console at https://console.aws.amazon.com/iam/.
- 3. In the left navigation pane, choose **Roles** and find the **amp-iamproxy-ingest-role** that you created in step 1.
- 4. Choose the **Trust relationships** tab and choose **Edit trust relationship**.
- 5. In the trust relationship policy JSON, replace aws-amp with adot-col and then choose **Update Trust Policy**. Your resulting trust policy should look like the following:

```
{
    "Version": "2012-10-17",
    "Statement": [
    {
        "Effect": "Allow",
        "Principal": {
            "Federated": "arn:aws:iam::account-id:oidc-provider/
oidc.eks.region.amazonaws.com/id/openid"
        },
        "Action": "sts:AssumeRoleWithWebIdentity",
        "Condition": {
            "StringEquals": {
                "oidc.eks.region.amazonaws.com/id/openid:sub":
            "system:serviceaccount:adot-col:amp-iamproxy-ingest-service-account"
        }
```

			}
		}	
]		
}			

6. Choose the **Permissions** tab and make sure that the following permissions policy is attached to the role.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
               "aps:RemoteWrite",
               "aps:GetSeries",
               "aps:GetLabels",
               "aps:GetMetricMetadata"
            ],
            "Resource": "*"
        }
    ]
}
```

Enabling Prometheus metric collection

1 Note

When you create a namespace in Amazon EKS, alertmanager and node exporter are disabled by default.

To enable Prometheus collection on an Amazon EKS or Kubernetes cluster

1. Fork and clone the sample app from the repository at <u>aws-otel-community</u>.

Then run the following commands.

```
cd ./sample-apps/prometheus-sample-app
docker build . -t prometheus-sample-app:latest
```

- 2. Push this image to a registry such as Amazon ECR or DockerHub.
- 3. Deploy the sample app in the cluster by copying this Kubernetes configuration and applying it. Change the image to the image that you just pushed by replacing {{PUBLIC_SAMPLE_APP_IMAGE}} in the prometheus-sample-app.yaml file.

```
curl https://raw.githubusercontent.com/aws-observability/aws-otel-collector/
main/examples/eks/aws-prometheus/prometheus-sample-app.yaml -o prometheus-sample-
app.yaml
kubectl apply -f prometheus-sample-app.yaml
```

4. Enter the following command to verify that the sample app has started. In the output of the command, you will see prometheus-sample-app in the NAME column.

kubectl get all -n aoc-prometheus-pipeline-demo

5. Start a default instance of the ADOT Collector. To do so, first enter the following command to pull the Kubernetes configuration for ADOT Collector.

curl https://raw.githubusercontent.com/aws-observability/aws-otel-collector/main/ examples/eks/aws-prometheus/prometheus-daemonset.yaml -o prometheus-daemonset.yaml

Then edit the template file, substituting the **remote_write** endpoint for your Amazon Managed Service for Prometheus workspace for YOUR_ENDPOINT and your Region for YOUR_REGION. Use the **remote_write** endpoint that is displayed in the Amazon Managed Service for Prometheus console when you look at your workspace details.

You'll also need to change YOUR_ACCOUNT_ID in the service account section of the Kubernetes configuration to your AWS account ID.

In this example, the ADOT Collector configuration uses an annotation (scrape=true) to tell which target endpoints to scrape. This allows the ADOT Collector to distinguish the sample app endpoint from kube-system endpoints in your cluster. You can remove this from the relabel configurations if you want to scrape a different sample app.

6. Enter the following command to deploy the ADOT collector.

kubectl apply -f prometheus-daemonset.yaml

7. Enter the following command to verify that the ADOT collector has started. Look for adotcol in the NAMESPACE column.

```
kubectl get pods -n adot-col
```

8. Verify that the pipeline works by using the logging exporter. Our example template is already integrated with the logging exporter. Enter the following commands.

```
kubectl get pods -A
kubectl logs -n adot-col name_of_your_adot_collector_pod
```

Some of the scraped metrics from the sample app will look like the following example.

```
Resource labels:
     -> service.name: STRING(kubernetes-service-endpoints)
     -> host.name: STRING(192.168.16.238)
     -> port: STRING(8080)
     -> scheme: STRING(http)
InstrumentationLibraryMetrics #0
Metric #0
Descriptor:
     -> Name: test_gauge0
     -> Description: This is my gauge
     -> Unit:
     -> DataType: DoubleGauge
DoubleDataPoints #0
StartTime: 0
Timestamp: 1606511460471000000
Value: 0.000000
```

9. To test whether Amazon Managed Service for Prometheus received the metrics, use awscurl. This tool enables you to send HTTP requests through the command line with AWS Sigv4 authentication, so you must have AWS credentials set up locally with the correct permissions to query from Amazon Managed Service for Prometheus For instructions on installing awscurl, see <u>awscurl</u>.

In the following command, replace AMP_REGION, and AMP_ENDPOINT with the information for your Amazon Managed Service for Prometheus workspace.

```
awscurl --service="aps" --region="AMP_REGION" "https://AMP_ENDPOINT/api/v1/query?
query=adot_test_gauge0"
{"status":"success","data":{"resultType":"vector","result":[{"metric":
{"___name___":"adot_test_gauge0"},"value":[1606512592.493,"16.87214000011479"]}]}}
```

If you receive a metric as the response, that means your pipeline setup has been successful and the metric has successfully propagated from the sample app into Amazon Managed Service for Prometheus.

Cleaning up

To clean up this demo, enter the following commands.

```
kubectl delete namespace aoc-prometheus-pipeline-demo
kubectl delete namespace adot-col
```

Advanced configuration

The Prometheus Receiver supports the full set of Prometheus scraping and re-labeling configurations described in <u>Configuration</u> in the Prometheus documentation. You can paste these configurations directly into your ADOT Collector configurations.

The configuration for the Prometheus Receiver includes your service discovery, scraping configurations, and re-labeling configurations. The receiver configuration looks like the following.

receivers:
 prometheus:
 config:
 [[Your Prometheus configuration]]

The following is an example configuration.

```
receivers:
prometheus:
config:
global:
    scrape_interval: 1m
    scrape_timeout: 10s
    scrape_configs:
    - job_name: kubernetes-service-endpoints
    sample_limit: 10000
    kubernetes_sd_configs:
        - role: endpoints
        tls_config:
            ca_file: /var/run/secrets/kubernetes.io/serviceaccount/ca.crt
```

```
insecure_skip_verify: true
bearer_token_file: /var/run/secrets/kubernetes.io/serviceaccount/token
```

If you have an existing Prometheus configuration, you must replace the \$ characters with \$\$ to avoid having the values replaced with environment variables. *This is especially important for the replacement value of the relabel_configurations. For example, if you start with the following relabel_configuration:

```
relabel_configs:
- source_labels:
[__meta_kubernetes_ingress_scheme, __address__, __meta_kubernetes_ingress_path]
regex: (.+);(.+);(.+)
replacement: ${1}://${2}${3}
target_label: __param_target
```

It would become the following:

```
relabel_configs:
- source_labels:
[__meta_kubernetes_ingress_scheme,__address__,__meta_kubernetes_ingress_path]
regex: (.+);(.+);(.+)
replacement: $${1}://${2}${3}
target_label: __param_target
```

Prometheus remote write exporter and Sigv4 authentication extension

The configuration for the Prometheus Remote Write Exporter and Sigv4 Authentication Extension are simpler than the Prometheus receiver. At this stage in the pipeline, metrics have already been ingested, and we're ready to export this data to Amazon Managed Service for Prometheus. The minimum requirement for a successful configuration to communicate with Amazon Managed Service for Prometheus is shown in the following example.

```
extensions:
    sigv4auth:
    service: "aps"
    region: "user-region"
exporters:
    prometheusremotewrite:
    endpoint: "https://aws-managed-prometheus-endpoint/api/v1/remote_write"
    auth:
        authenticator: "sigv4auth"
```

This configuration sends an HTTPS request that is signed by AWS SigV4 using AWS credentials from the default AWS credentials chain. For more information, see <u>Configuring the AWS SDK for</u> Go. You must specify the service to be aps.

Regardless of the method of deployment, the ADOT collector must have access to one of the listed options in the default AWS credentials chain. The Sigv4 Authentication Extension depends on the AWS SDK for Go and uses it to fetch credentials and authenticate. You must ensure that these credentials have remote write permissions for Amazon Managed Service for Prometheus.

Set up metrics ingestion from Amazon ECS using AWS Distro for Open Telemetry

This section explains how to collect metrics from Amazon Elastic Container Service (Amazon ECS) and ingest them into Amazon Managed Service for Prometheus using AWS Distro for Open Telemetry (ADOT). It also describes how to visualize your metrics in Amazon Managed Grafana.

Prerequisites

🛕 Important

Before you begin, you must have an Amazon ECS environment on an AWS Fargate cluster with default settings, an Amazon Managed Service for Prometheus workspace, and an Amazon Managed Grafana workspace. We assume that you are familiar with container workloads, Amazon Managed Service for Prometheus, and Amazon Managed Grafana.

For more information, see the following links:

- For information about how to create an Amazon ECS environment on a Fargate cluster with default settings, see <u>Creating a cluster</u> in the *Amazon ECS Developer Guide*.
- For information about how to create an Amazon Managed Service for Prometheus workspace, see <u>Create a workspace</u> in the Amazon Managed Service for Prometheus User Guide.
- For information about how to create an Amazon Managed Grafana workspace, see <u>Creating a</u> workspace in the Amazon Managed Grafana User Guide.

Step 1: Define a custom ADOT collector container image

Use the following config file as a template to define your own ADOT collector container image. Replace *my*-*remote*-*URL* and *my*-*region* with your endpoint and region values. Save the config in a file called *adot*-*config.yaml*.

User Guide

(i) Note

This configuration uses the sigv4auth extension to authenticate calls to Amazon Managed Service for Prometheus. For more information about configuring sigv4auth, see <u>Authenticator - Sigv4</u> on GitHub.

```
receivers:
  prometheus:
    config:
      global:
        scrape_interval: 15s
        scrape_timeout: 10s
      scrape_configs:
      - job_name: "prometheus"
        static_configs:
        - targets: [ 0.0.0.0:9090 ]
  awsecscontainermetrics:
    collection_interval: 10s
processors:
  filter:
    metrics:
      include:
        match_type: strict
        metric_names:
          - ecs.task.memory.utilized
          - ecs.task.memory.reserved
          - ecs.task.cpu.utilized
          - ecs.task.cpu.reserved
          - ecs.task.network.rate.rx
          - ecs.task.network.rate.tx
          - ecs.task.storage.read_bytes
          - ecs.task.storage.write_bytes
exporters:
  prometheusremotewrite:
    endpoint: my-remote-URL
    auth:
      authenticator: sigv4auth
  logging:
    loglevel: info
extensions:
  health_check:
```

```
pprof:
    endpoint: :1888
  zpages:
    endpoint: :55679
  sigv4auth:
    region: my-region
    service: aps
service:
  extensions: [pprof, zpages, health_check, sigv4auth]
  pipelines:
    metrics:
      receivers: [prometheus]
      exporters: [logging, prometheusremotewrite]
    metrics/ecs:
      receivers: [awsecscontainermetrics]
      processors: [filter]
      exporters: [logging, prometheusremotewrite]
```

Step 2: Push your ADOT collector container image to an Amazon ECR repository

Use a Dockerfile to create and push your container image to an Amazon Elastic Container Registry (ECR) repository.

1. Build the Dockerfile to copy and add your container image to the OTEL Docker image.

```
FROM public.ecr.aws/aws-observability/aws-otel-collector:latest
COPY adot-config.yaml /etc/ecs/otel-config.yaml
CMD ["--config=/etc/ecs/otel-config.yaml"]
```

2. Create an Amazon ECR repository.

3. Create your container image.

```
# build ADOT collector image:
docker build -t $COLLECTOR_REPOSITORY:ecs .
```

í) Note

This assumes you are building your container in the same environment that it will run in. If not, you may need to use the --platform parameter when building the image.

4. Sign in to the Amazon ECR repository. Replace *my*-*region* with your region value.

```
# sign in to repo:
aws ecr get-login-password --region my-region | \
docker login --username AWS --password-stdin $COLLECTOR_REPOSITORY
```

5. Push your container image.

```
# push ADOT collector image:
docker push $COLLECTOR_REPOSITORY:ecs
```

Step 3: Create an Amazon ECS task definition to scrape Amazon Managed Service for Prometheus

Create an Amazon ECS task definition to scrape Amazon Managed Service for Prometheus. Your task definition should include a container named adot-collector and a container named prometheus. prometheus generates metrics, and adot-collector scrapes prometheus.

i Note

Amazon Managed Service for Prometheus runs as a service, collecting metrics from containers. The containers in this case run Prometheus locally, in Agent mode, which send the local metrics to Amazon Managed Service for Prometheus.

Example: Task definition

The following is an example of how your task definition might look. You can use this example as a template to create your own task definition. Replace the image value of adot-collector with your repository URL and image tag (\$COLLECTOR_REPOSITORY:ecs). Replace the region values of adot-collector and prometheus with your region values.

```
"family": "adot-prom",
  "networkMode": "awsvpc",
  "containerDefinitions": [
    {
      "name": "adot-collector",
      "image": "account_id.dkr.ecr.region.amazonaws.com/image-tag",
      "essential": true,
      "logConfiguration": {
        "logDriver": "awslogs",
        "options": {
          "awslogs-group": "/ecs/ecs-adot-collector",
          "awslogs-region": "my-region",
          "awslogs-stream-prefix": "ecs",
          "awslogs-create-group": "True"
        }
      }
    },
    {
      "name": "prometheus",
      "image": "prom/prometheus:main",
      "logConfiguration": {
        "logDriver": "awslogs",
        "options": {
          "awslogs-group": "/ecs/ecs-prom",
          "awslogs-region": "my-region",
          "awslogs-stream-prefix": "ecs",
          "awslogs-create-group": "True"
        }
      }
    }
  ],
  "requiresCompatibilities": [
    "FARGATE"
  ],
  "cpu": "1024"
}
```

Step 4: Give your task permissions to access Amazon Managed Service for Prometheus

To send the scraped metrics to Amazon Managed Service for Prometheus, your Amazon ECS task must have the correct permissions to call the AWS API operations for you. You must create an IAM role for your tasks and attach the AmazonPrometheusRemoteWriteAccess policy to it. For more

information about creating this role and attaching the policy, see <u>Creating an IAM role and policy</u> for your tasks.

After you attach AmazonPrometheusRemoteWriteAccess to your IAM role, and use that role for your tasks, Amazon ECS can send your scraped metrics to Amazon Managed Service for Prometheus.

Step 5: Visualize your metrics in Amazon Managed Grafana

🛕 Important

Before you begin, you must run a Fargate task on your Amazon ECS task definition. Otherwise, Amazon Managed Service for Prometheus can't consume your metrics.

- 1. From the navigation pane in your Amazon Managed Grafana workspace, choose **Data sources** under the AWS icon.
- 2. On the **Data sources** tab, for **Service**, select **Amazon Managed Service for Prometheus** and choose your **Default Region**.
- 3. Choose Add data source.
- 4. Use the ecs and prometheus prefixes to query and view your metrics.

Set up metrics ingestion from an Amazon EC2 instance using remote write

This section explains how to run a Prometheus server with remote write in an Amazon Elastic Compute Cloud (Amazon EC2) instance. It explains how to collect metrics from a demo application written in Go and send them to an Amazon Managed Service for Prometheus workspace.

Prerequisites

<u> Important</u>

Before you start, you must have installed Prometheus v2.26 or later. We assume that you're familiar with Prometheus, Amazon EC2, and Amazon Managed Service for Prometheus. For information about how to install Prometheus, see <u>Getting started</u> on the Prometheus website.

If you're unfamiliar with Amazon EC2 or Amazon Managed Service for Prometheus, we recommend that you start by reading the following sections:

- What is Amazon Elastic Compute Cloud?
- What is Amazon Managed Service for Prometheus?

Create an IAM role for Amazon EC2

To stream metrics, you must first create an IAM role with the AWS managed policy **AmazonPrometheusRemoteWriteAccess**. Then, you can launch an instance with the role and stream metrics into your Amazon Managed Service for Prometheus workspace.

- 1. Open the IAM console at https://console.aws.amazon.com/iam/.
- 2. From the navigation pane, choose **Roles**, and then choose **Create role**.
- For the type of trusted entity, choose AWS service. For the use case, choose EC2. Choose Next: Permissions.
- 4. In the search bar, enter AmazonPrometheusRemoteWriteAccess. For Policy name, select AmazonPrometheusRemoteWriteAccess, and then choose Attach policy. Choose Next:Tags.
- 5. (Optional) Create IAM tags for your IAM role. Choose **Next: Review**.
- 6. Enter a name for your role. Choose **Create policy**.

Launch an Amazon EC2 instance

To launch an Amazon EC2 instance, follow the instructions at <u>Launch an instance</u> in the Amazon Elastic Compute Cloud User Guide for Linux Instances.

Run the demo application

After creating your IAM role, and launching an EC2 instance with the role, you can run a demo application to see it work.

To run a demo application and test metrics

1. Use the following template to create a Go file named main.go.

package main import (

```
"github.com/prometheus/client_golang/prometheus/promhttp"
    "net/http"
)
func main() {
    http.Handle("/metrics", promhttp.Handler())
    http.ListenAndServe(":8000", nil)
}
```

2. Run the following commands to install the correct dependencies.

```
sudo yum update -y
sudo yum install -y golang
go get github.com/prometheus/client_golang/prometheus/promhttp
```

3. Run the demo application.

```
go run main.go
```

The demo application should run on port 8000 and show all of the exposed Prometheus metrics. The following is an example of these metrics.

```
curl -s http://localhost:8000/metrics
. . .
process_max_fds 4096# HELP process_open_fds Number of open file descriptors.# TYPE
process_open_fds gauge
process_open_fds 10# HELP process_resident_memory_bytes Resident memory size in
bytes.# TYPE process_resident_memory_bytes gauge
process_resident_memory_bytes 1.0657792e+07# HELP process_start_time_seconds Start
time of the process since unix epoch in seconds.# TYPE process_start_time_seconds
gauge
process_start_time_seconds 1.61131955899e+09# HELP process_virtual_memory_bytes
Virtual memory size in bytes.# TYPE process_virtual_memory_bytes gauge
process_virtual_memory_bytes 7.77281536e+08# HELP process_virtual_memory_max_bytes
Maximum amount of virtual memory available in bytes.# TYPE
process_virtual_memory_max_bytes gauge
process_virtual_memory_max_bytes -1# HELP
promhttp_metric_handler_requests_in_flight Current number of scrapes being
 served.# TYPE promhttp_metric_handler_requests_in_flight gauge
```

```
promhttp_metric_handler_requests_in_flight 1# HELP
promhttp_metric_handler_requests_total Total number of scrapes by HTTP status
code.# TYPE promhttp_metric_handler_requests_total counter
promhttp_metric_handler_requests_total{code="200"} 1
promhttp_metric_handler_requests_total{code="500"} 0
promhttp_metric_handler_requests_total{code="503"} 0
```

Create an Amazon Managed Service for Prometheus workspace

To create an Amazon Managed Service for Prometheus workspace, follow the instructions at <u>Create</u> <u>a workspace</u>.

Run a Prometheus server

 Use the following example YAML file as a template to create a new file named prometheus.yaml. For url, replace my-region with your Region value and myworkspace-id with the workspace ID that Amazon Managed Service for Prometheus generated for you. For region, replace my-region with your Region value.

Example: YAML file

```
global:
 scrape_interval: 15s
  external_labels:
    monitor: 'prometheus'
scrape_configs:
  - job_name: 'prometheus'
    static_configs:
      - targets: ['localhost:8000']
remote_write:
    url: https://aps-workspaces.my-region.amazonaws.com/workspaces/my-workspace-id/
api/v1/remote_write
    queue_config:
        max_samples_per_send: 1000
        max_shards: 200
        capacity: 2500
    sigv4:
         region: my-region
```

2. Run the Prometheus server to send the demo application's metrics to your Amazon Managed Service for Prometheus workspace.

```
prometheus --config.file=prometheus.yaml
```

The Prometheus server should now send the demo application's metrics to your Amazon Managed Service for Prometheus workspace.

Using a Prometheus instance as a collector

You can use a Prometheus instance, running in *agent* mode (known as a *Prometheus agent*), to scrape metrics and send them to your Amazon Managed Service for Prometheus workspace.

The following topics describe different ways to set up a Prometheus instance running in agent mode as a collector for your metrics.

🔥 Warning

When you create a Prometheus agent, you are responsible for its configuration and maintenance. Avoid exposing Prometheus scrape endpoints to the public internet by enabling security features.

If you set up multiple Prometheus instances that monitor the same set of metrics and sent them to a single Amazon Managed Service for Prometheus workspace for high availability, you need to set up deduplication. If you don't follow the steps to set up deduplication, you will be charged for all data samples sent to Amazon Managed Service for Prometheus, including duplicate samples. For instructions about setting up deduplication, see <u>Deduplicating high availability metrics sent to Amazon Managed Service for Prometheus</u>.

Topics

- Set up ingestion from a new Prometheus server using Helm
- Set up ingestion from an existing Prometheus server in Kubernetes on EC2
- Set up ingestion from an existing Prometheus server in Kubernetes on Fargate

Set up ingestion from a new Prometheus server using Helm

The instructions in this section get you up and running with Amazon Managed Service for Prometheus quickly. You set up a new Prometheus server in an Amazon EKS cluster, and the new server uses a default configuration to send metrics to Amazon Managed Service for Prometheus. This method has the following prerequisites:

- You must have an Amazon EKS cluster from which the new Prometheus server will collect metrics.
- Your Amazon EKS cluster must have an Amazon EBS CSI driver installed (required by Helm).
- You must use Helm CLI 3.0 or later.
- You must use a Linux or macOS computer to perform the steps in the following sections.

Step 1: Add new Helm chart repositories

To add new Helm chart repositories, enter the following commands. For more information about these commands, see <u>Helm Repo</u>.

helm repo add prometheus-community https://prometheus-community.github.io/helm-charts helm repo add kube-state-metrics https://kubernetes.github.io/kube-state-metrics helm repo update

Step 2: Create a Prometheus namespace

Enter the following command to create a Prometheus namespace for the Prometheus server and other monitoring components. Replace *prometheus-namespace* with the name that you want for this namespace.

```
kubectl create namespace prometheus-namespace
```

Step 3: Set up IAM roles for service accounts

For the method of onboarding that we are documenting, you need to use IAM roles for service accounts in the Amazon EKS cluster where the Prometheus server is running.

With IAM roles for service accounts, you can associate an IAM role with a Kubernetes service account. This service account can then provide AWS permissions to the containers in any pod that uses that service account. For more information, see IAM roles for service accounts.

If you have not already set up these roles, follow the instructions at <u>Set up service roles for</u> <u>the ingestion of metrics from Amazon EKS clusters</u> to set up the roles. The instructions in that section require the use of eksctl. For more information, see <u>Getting started with Amazon Elastic</u> Kubernetes Service – eksctl.

Note

When you are not on EKS or AWS and using just access key and secret key to access Amazon Managed Service for Prometheus, you cannot use the EKS-IAM-ROLE based SigV4.

Step 4: Set up the new server and start ingesting metrics

To install the new Prometheus server that sends metrics to your Amazon Managed Service for Prometheus workspace, follow these steps.

To install a new Prometheus server to send metrics to your Amazon Managed Service for Prometheus workspace

- Use a text editor to create a file named my_prometheus_values_yaml with the following content.
 - Replace IAM_PROXY_PROMETHEUS_ROLE_ARN with the ARN of the amp-iamproxy-ingestrole that you created in <u>Set up service roles for the ingestion of metrics from Amazon EKS</u> clusters.
 - Replace WORKSPACE_ID with the ID of your Amazon Managed Service for Prometheus workspace.
 - Replace *REGION* with the Region of your Amazon Managed Service for Prometheus workspace.

```
## The following is a set of default values for prometheus server helm chart which
enable remoteWrite to AMP
## For the rest of prometheus helm chart values see: https://github.com/prometheus-
community/helm-charts/blob/main/charts/prometheus/values.yaml
##
serviceAccounts:
    server:
    name: amp-iamproxy-ingest-service-account
    annotations:
```

```
eks.amazonaws.com/role-arn: ${IAM_PROXY_PROMETHEUS_ROLE_ARN}
server:
    remoteWrite:
        url: https://aps-workspaces.${REGION}.amazonaws.com/workspaces/
${WORKSPACE_ID}/api/v1/remote_write
        sigv4:
        region: ${REGION}
        queue_config:
        max_samples_per_send: 1000
        max_shards: 200
        capacity: 2500
```

- 2. Enter the following command to create the Prometheus server.
 - Replace *prometheus-chart-name* with your Prometheus release name.
 - Replace *prometheus-namespace* with the name of your Prometheus namespace.

```
helm install prometheus-chart-name prometheus-community/prometheus -n prometheus-
namespace \
```

-f my_prometheus_values_yaml

🚯 Note

You can customize the helm install command in many ways. For more information, see Helm install in the Helm documentation.

Set up ingestion from an existing Prometheus server in Kubernetes on EC2

Amazon Managed Service for Prometheus supports ingesting metrics from Prometheus servers in clusters running Amazon EKS and in self-managed Kubernetes clusters running on Amazon EC2. The detailed instructions in this section are for a Prometheus server in an Amazon EKS cluster. The steps for a self-managed Kubernetes cluster on Amazon EC2 are the same, except that you will need to set up the OIDC provider and IAM roles for service accounts yourself in the Kubernetes cluster.

The instructions in this section use Helm as the Kubernetes package manager.

Topics

- Step 1: Set up IAM roles for service accounts
- Step 2: Upgrade your existing Prometheus server using Helm

Step 1: Set up IAM roles for service accounts

For the method of onboarding that we are documenting, you need to use IAM roles for service accounts in the Amazon EKS cluster where the Prometheus server is running. These roles are also called *service roles*.

With service roles, you can associate an IAM role with a Kubernetes service account. This service account can then provide AWS permissions to the containers in any pod that uses that service account. For more information, see IAM roles for service accounts.

If you have not already set up these roles, follow the instructions at <u>Set up service roles for the</u> ingestion of metrics from Amazon EKS clusters to set up the roles.

Step 2: Upgrade your existing Prometheus server using Helm

The instructions in this section include setting up remote write and sigv4 to authenticate and authorize the Prometheus server to remote write to your Amazon Managed Service for Prometheus workspace.

Using Prometheus version 2.26.0 or later

Follow these steps if you are using a Helm chart with Prometheus Server image of version 2.26.0 or later.

To set up remote write from a Prometheus server using Helm chart

- 1. Create a new remote write section in your Helm configuration file:
 - Replace \${IAM_PROXY_PROMETHEUS_ROLE_ARN} with the ARN of the amp-iamproxyingest-role that you created in <u>Step 1: Set up IAM roles for service accounts</u>. The role ARN should have the format of arn:aws:iam::your account ID:role/amp-iamproxyingest-role.
 - Replace \${WORKSPACE_ID} with your Amazon Managed Service for Prometheus workspace ID.
 - Replace \${REGION} with the Region of the Amazon Managed Service for Prometheus workspace (such as us-west-2).

```
## The following is a set of default values for prometheus server helm chart which
enable remoteWrite to AMP
    ## For the rest of prometheus helm chart values see: https://github.com/
prometheus-community/helm-charts/blob/main/charts/prometheus/values.yaml
    ##
    serviceAccounts:
      server:
        name: amp-iamproxy-ingest-service-account
        annotations:
          eks.amazonaws.com/role-arn: ${IAM_PROXY_PROMETHEUS_ROLE_ARN}
    server:
      remoteWrite:
        - url: https://aps-workspaces.${REGION}.amazonaws.com/workspaces/
${WORKSPACE_ID}/api/v1/remote_write
          sigv4:
            region: ${REGION}
          queue_config:
            max_samples_per_send: 1000
            max_shards: 200
            capacity: 2500
```

- 2. Update your existing Prometheus Server configuration using Helm:
 - Replace prometheus-chart-name with your Prometheus release name.
 - Replace prometheus-namespace with the Kubernetes namespace where your Prometheus Server is installed.
 - Replace my_prometheus_values_yaml with the path to your Helm configuration file.
 - Replace current_helm_chart_version with the current version of your Prometheus Server Helm chart. You can find the current chart version by using the helm list command.

```
helm upgrade prometheus-chart-name prometheus-community/prometheus \
    -n prometheus-namespace \
    -f my_prometheus_values_yaml \
    --version current_helm_chart_version
```

Using earlier versions of Prometheus

Follow these steps if you are using a version of Prometheus earlier than 2.26.0. These steps use a sidecar approach, because earlier versions of Prometheus don't natively support AWS Signature Version 4 signing process (AWS SigV4).

These instructions assume that you are using Helm to deploy Prometheus.

To set up remote write from a Prometheus server

1. On your Prometheus server, create a new remote write configuration. First, create a new update file. We will call the file amp_ingest_override_values.yaml.

Add the following values to the YAML file.

```
serviceAccounts:
        server:
            name: "amp-iamproxy-ingest-service-account"
            annotations:
                eks.amazonaws.com/role-arn:
 "${SERVICE_ACCOUNT_IAM_INGEST_ROLE_ARN}"
    server:
        sidecarContainers:
            - name: aws-sigv4-proxy-sidecar
              image: public.ecr.aws/aws-observability/aws-sigv4-proxy:1.0
              args:
              - --name
              - aps
              - --region
              - ${REGION}
              - --host
              - aps-workspaces.${REGION}.amazonaws.com
              - --port
              - :8005
              ports:
              - name: aws-sigv4-proxy
                containerPort: 8005
        statefulSet:
            enabled: "true"
        remoteWrite:
            - url: http://localhost:8005/workspaces/${WORKSPACE_ID}/api/v1/
remote_write
```

Replace \${REGION} with the Region of the Amazon Managed Service for Prometheus workspace.

Replace \${SERVICE_ACCOUNT_IAM_INGEST_ROLE_ARN} with the ARN of the **amp-iamproxy-ingest-role** that you created in <u>Step 1: Set up IAM roles for service accounts</u>. The role ARN should have the format of arn:aws:iam::your account ID:role/amp-iamproxy-ingest-role.

Replace \${WORKSPACE_ID} with your workspace ID.

2. Upgrade your Prometheus Helm chart. First, find your Helm chart name by entering the following command. In the output from this command, look for a chart with a name that includes prometheus.

helm ls --all-namespaces

Then enter the following command.

```
helm upgrade --install prometheus-helm-chart-name prometheus-community/prometheus -
n prometheus-namespace -f ./amp_ingest_override_values.yaml
```

Replace *prometheus-helm-chart-name* with the name of the Prometheus helm chart returned in the previous command. Replace *prometheus-namespace* with the name of your namespace.

Downloading Helm charts

If you don't already have Helm charts downloaded locally, you can use the following command to download them.

```
helm repo add prometheus-community https://prometheus-community.github.io/helm-charts
    helm pull prometheus-community/prometheus --untar
```

Set up ingestion from an existing Prometheus server in Kubernetes on Fargate

Amazon Managed Service for Prometheus supports ingesting metrics from Prometheus servers in self-managed Kubernetes clusters running on Fargate. To ingest metrics from Prometheus servers

in Amazon EKS clusters running on Fargate, override the default configs in a config file named amp_ingest_override_values.yaml as follows:

```
prometheus-node-exporter:
        enabled: false
    alertmanager:
        enabled: false
    serviceAccounts:
      server:
        name: amp-iamproxy-ingest-service-account
        annotations:
          eks.amazonaws.com/role-arn: ${IAM_PROXY_PROMETHEUS_ROLE_ARN}
    server:
      persistentVolume:
        enabled: false
      remoteWrite:
        - url: https://aps-workspaces.${REGION}.amazonaws.com/workspaces/
${WORKSPACE_ID}/api/v1/remote_write
          sigv4:
            region: ${REGION}
          queue_config:
            max_samples_per_send: 1000
            max_shards: 200
            capacity: 2500
```

Install Prometheus using the overrides with the following command:

Note that in the Helm chart configuration we disabled the node exporter and the alert manager as well as running the Prometheus server deployment.

You can verify the install with the following example test query.

```
$ awscurl --region region --service aps "https://aps-
workspaces.region_id.amazonaws.com/workspaces/workspace_id/api/v1/query?
query=prometheus_api_remote_read_queries"
```

```
{"status":"success","data":{"resultType":"vector","result":[{"metric":
{"__name__":"prometheus_api_remote_read_queries","instance":"localhost:9090","job":"prometheus"
[1648461236.419,"0"]}]}21
```

Set up Amazon Managed Service for Prometheus for high availability data

When you send data to Amazon Managed Service for Prometheus, it is automatically replicated across AWS Availability Zones in the Region, and is served to you from a cluster of hosts that provide scalability, availability, and security. You might want to add additional high availability fail-safes, depending on your particular setup. There are two common ways that you might additional high availability safeties to your setup:

 If you have multiple containers or instances that have the same data, you can send that data to Amazon Managed Service for Prometheus and have the data automatically de-duplicated. This helps to ensure that your data will be sent to your Amazon Managed Service for Prometheus workspace.

For more information about de-duplicating high-availability data, see <u>Deduplicating high</u> availability metrics sent to Amazon Managed Service for Prometheus.

• If you want to ensure that you have access to your data, even when the AWS Region is not available, you can send your metrics to a second workspace, in another Region.

For more information about sending metrics data to multiple workspaces, see <u>Use cross Region</u> workspaces to add high availability in Amazon Managed Service for Prometheus.

Topics

- Deduplicating high availability metrics sent to Amazon Managed Service for Prometheus
- Send high availability data to Amazon Managed Service for Prometheus with Prometheus
- Set up high availability data to Amazon Managed Service for Prometheus using the Prometheus Operator Helm chart
- Send high-availability data to Amazon Managed Service for Prometheus with AWS Distro for OpenTelemetry
- Send high availability data to Amazon Managed Service for Prometheus with the Prometheus community Helm chart

- Answers to common questions about high availability configuration in Amazon Managed Service for Prometheus
- <u>Use cross Region workspaces to add high availability in Amazon Managed Service for</u> Prometheus

Deduplicating high availability metrics sent to Amazon Managed Service for Prometheus

You can send data from multiple Prometheus *agents* (Prometheus instances running in Agent mode) to your Amazon Managed Service for Prometheus workspace. If some of these instances are recording and sending the same metrics, your data will have a higher availability (even if one of the agents stops sending data, the Amazon Managed Service for Prometheus workspace will still receive the data from another instance). However, you want your Amazon Managed Service for Prometheus workspace to automatically de-duplicate the metrics so that you don't see the metrics multiple times, and aren't charged for the data ingestion and storage multiple times.

For Amazon Managed Service for Prometheus to automatically de-duplicate data from multiple Prometheus agents, you give the set of agents that are sending the duplicate data a single *cluster name*, and each of the instances a *replica name*. The cluster name identifies the instances as having shared data, and the replica name allows Amazon Managed Service for Prometheus to identify the source of each metric. The final stored metrics include the cluster label, but not the replica, so the metrics appear to be coming from a single source.

1 Note

Certain versions of Kubernetes (1.28 and 1.29) may emit their own metric with a cluster label. This can cause issues with Amazon Managed Service for Prometheus deduplication. See the High availability FAQ for more information.

The following topics show how to send data and include the cluster and __replica__ labels, so that Amazon Managed Service for Prometheus de-duplicates the data automatically.

🔥 Important

If you do not set up deduplication, you will be charged for all data samples that are sent to Amazon Managed Service for Prometheus. These data samples include duplicate samples.

Send high availability data to Amazon Managed Service for Prometheus with Prometheus

To set up a high availability configuration with Prometheus, you must apply external labels on all instances of a high availability group, so Amazon Managed Service for Prometheus can identify them. Use the cluster label to identify a Prometheus instance agent as part of a high availability group. Use the __replica__ label to identify each replica in the group separately. You need to apply both __replica__ and cluster labels for de-duplication to work.

i Note

The <u>replica</u> label is formatted with two underscore symbols before and after the word replica.

Example: code snippets

In the following code snippets, the cluster label identifies the Prometheus instance agent promteam1, and the _replica_ label identifies the replicas replica1 and replica2.

```
cluster: prom-team1
__replica__: replica1
```

```
cluster: prom-team1
__replica_: replica2
```

As Amazon Managed Service for Prometheus stores data samples from high availability replicas with these labels, it strips the replica label when the samples are accepted. This means that you will only have a 1:1 series mapping for your current series instead of a series per replica. The cluster label is kept.

🚯 Note

Certain versions of Kubernetes (1.28 and 1.29) may emit their own metric with a cluster label. This can cause issues with Amazon Managed Service for Prometheus deduplication. See the High availability FAQ for more information.

Set up high availability data to Amazon Managed Service for Prometheus using the Prometheus Operator Helm chart

To set up a high availability configuration with the Prometheus Operator in Helm, you must apply external labels on all instances of a high availability group, so Amazon Managed Service for Prometheus can identify them. You also must set the attributes replicaExternalLabelName and externalLabels on the Prometheus Operator Helm chart.

Example: YAML header

In the following YAML header, cluster is added to externalLabel to identify a Prometheus instance agent as part of a high-availability group, and replicaExternalLabels identifies each replica in the group.

```
replicaExternalLabelName: __replica__
externalLabels:
cluster: prom-dev
```

🚯 Note

Certain versions of Kubernetes (1.28 and 1.29) may emit their own metric with a cluster label. This can cause issues with Amazon Managed Service for Prometheus deduplication. See the High availability FAQ for more information.

Send high-availability data to Amazon Managed Service for Prometheus with AWS Distro for OpenTelemetry

AWS Distro for OpenTelemetry (ADOT) is a secure and production-ready distribution of the OpenTelemetry project. ADOT provides you with source APIs, libraries, and agents, so you can collect distributed traces and metrics for application monitoring. For information about ADOT, see <u>About AWS Distro for Open Telemetry</u>.

To set up ADOT with a high availability configuration, you must configure an ADOT collector container image and apply the external labels cluster and __replica__ to the AWS Prometheus remote write exporter. This exporter sends your scraped metrics to your Amazon Managed Service for Prometheus workspace via the remote_write endpoint. When you set these labels on the remote write exporter, you prevent duplicate metrics from being kept while redundant replicas run. For more information about the AWS Prometheus remote write exporter,

see <u>Getting started with Prometheus remote write exporter for Amazon Managed Service for</u> Prometheus.

i Note

Certain versions of Kubernetes (1.28 and 1.29) may emit their own metric with a cluster label. This can cause issues with Amazon Managed Service for Prometheus deduplication. See the <u>High availability FAQ</u> for more information.

Send high availability data to Amazon Managed Service for Prometheus with the Prometheus community Helm chart

To set up a high availability configuration with the Prometheus community Helm chart, you must apply external labels on all instances of a high availability group, so Amazon Managed Service for Prometheus can identify them. Here is an example of how you could add the external_labels to a single instance of Prometheus from the Prometheus community Helm chart.

```
server:
global:
    external_labels:
        cluster: monitoring-cluster
        __replica__: replica-1
```

🚺 Note

If you want multiple replicas, you have to deploy the chart multiple times with different replica values, because the Prometheus community Helm chart does not let you dynamically set the replica value when increasing the number of replicas directly from the controller group. If you prefer to have the replica label auto-set, use the prometheus-operator Helm chart.

(i) Note

Certain versions of Kubernetes (1.28 and 1.29) may emit their own metric with a cluster label. This can cause issues with Amazon Managed Service for Prometheus deduplication. See the High availability FAQ for more information.

Answers to common questions about high availability configuration in Amazon Managed Service for Prometheus

Should I include the value __replica__ into another label to track the sample points?

In a high availability setting, Amazon Managed Service for Prometheus ensures data samples are not duplicated by electing a leader in the cluster of Prometheus instances. If the leader replica stops sending data samples for 30 seconds, Amazon Managed Service for Prometheus automatically makes another Prometheus instance a leader replica and ingests data from the new leader, including any missed data. Therefore, the answer is no, it is not recommended. Doing so may cause issues like:

- Querying a count in **PromQL** may return higher than expected value during the period of electing a new leader.
- The number of active series gets increased during a period of electing a new leader and it reaches the active series limits. See AMP Quotas for more info.

Kubernetes seems to have it's own *cluster* label, and is not deduplicating my metrics. How can I fix this?

A new metric, apiserver_storage_size_bytes was introduced in Kubernetes 1.28, with a cluster label. This can cause issues with deduplication in Amazon Managed Service for Prometheus, which depends on the cluster label. In Kubernetes 1.3, the label is renamed to storage-cluster_id (it is also renamed in later patches of 1.28 and 1.29). If your cluster is emitting this metric with the cluster label, Amazon Managed Service for Prometheus can't dedupe the associated time series. We recommend you upgrade your Kubernetes cluster to the latest patched version to avoid this problem. Alternately, you can relabel the cluster label on your apiserver_storage_size_bytes metric before ingesting it into Amazon Managed Service for Prometheus.

Note

For more details about the change to Kubernetes, see <u>Rename Label cluster to</u> storage_cluster_id for apiserver_storage_size_bytes metric in the *Kubernetes GitHub project*.

Use cross Region workspaces to add high availability in Amazon Managed Service for Prometheus

To add cross-Region availability to your data, you can send metrics to multiple workspaces across AWS Regions. Prometheus supports both multiple writers and cross-Region writing.

The following example shows how to set up a Prometheus server running in Agent mode to send metrics to two workspaces in different Regions with Helm.

```
extensions:
      sigv4auth:
        service: "aps"
    receivers:
      prometheus:
        config:
          scrape_configs:
            - job_name: 'kubernetes-kubelet'
              scheme: https
              tls_config:
                ca_file: /var/run/secrets/kubernetes.io/serviceaccount/ca.crt
                insecure_skip_verify: true
              bearer_token_file: /var/run/secrets/kubernetes.io/serviceaccount/token
              kubernetes_sd_configs:
              - role: node
              relabel_configs:
              - action: labelmap
                regex: __meta_kubernetes_node_label_(.+)
              - target_label: __address__
                replacement: kubernetes.default.svc.cluster.local:443
              - source_labels: [___meta_kubernetes_node_name]
                regex: (.+)
                target_label: __metrics_path__
                replacement: /api/v1/nodes/$${1}/proxy/metrics
    exporters:
      prometheusremotewrite/one:
        endpoint: "https://aps-workspaces.workspace_1_region.amazonaws.com/workspaces/
ws-workspace_1_id/api/v1/remote_write"
        auth:
          authenticator: sigv4auth
      prometheusremotewrite/two:
```

```
endpoint: "https://aps-workspaces.workspace_2_region.amazonaws.com/workspaces/
ws-workspace_2_id/api/v1/remote_write"
    auth:
    auth:
    authenticator: sigv4auth
service:
    extensions: [sigv4auth]
    pipelines:
    metrics/one:
    receivers: [prometheus]
    exporters: [prometheusremotewrite/one]
    metrics/two:
    receivers: [prometheus]
    exporters: [prometheus]
```

Query your Prometheus metrics

Now that metrics are being ingested to the workspace, you can query them.

To create dashboards with visual representations of your metrics, you can use a service such as Amazon Managed Grafana. Amazon Managed Grafana (or a standalone instance of Grafana) can build a graphical interface that shows your metrics in a wide variety of display presentation styles. For more information about Amazon Managed Grafana see the <u>Amazon Managed Grafana User</u> <u>Guide</u>.

You can also create one-off queries, explore your data, or write your own applications that use your metrics by using direct queries. Direct queries use the Amazon Managed Service for Prometheus API and the standard Prometheus query language, PromQL, to get data from your Prometheus workspace. For more information about PromQL and its syntax, see <u>Querying Prometheus</u> in the Prometheus documentation.

Topics

- Secure your metric queries
- Set up Amazon Managed Grafana for use with Amazon Managed Service for Prometheus
- <u>Set up Grafana open source or Grafana Enterprise for use with Amazon Managed Service for</u>
 <u>Prometheus</u>
- Query using Grafana running in an Amazon EKS cluster
- Query using Prometheus-compatible APIs
- Get statistics about your query usage for each query

Secure your metric queries

Amazon Managed Service for Prometheus provides ways of helping you secure the querying of your metrics.

Using AWS PrivateLink with Amazon Managed Service for Prometheus

The network traffic for querying metrics in Amazon Managed Service for Prometheus can be done over a public internet endpoint, or by a VPC endpoint through AWS PrivateLink. When you use AWS PrivateLink, network traffic from your VPCs is secured within the AWS network without going over the public internet. To create an AWS PrivateLink VPC endpoint for Amazon Managed Service for Prometheus, see Using Amazon Managed Service for Prometheus with interface VPC endpoints.

Authentication and authorization

AWS Identity and Access Management is a web service that helps you securely control access to AWS resources. You use IAM to control who is authenticated (signed in) and authorized (has permissions) to use resources. Amazon Managed Service for Prometheus integrates with IAM to help you keep your data secure. When you set up Amazon Managed Service for Prometheus, you'll need to create some IAM roles that enable Grafana servers to query metrics stored in Amazon Managed Service for Prometheus workspaces. For more information about IAM, see <u>What is IAM?</u>.

Another AWS security feature that can help you set up Amazon Managed Service for Prometheus is the AWS Signature Version 4 signing process (AWS SigV4). Signature Version 4 is the process to add authentication information to AWS requests sent by HTTP. For security, most requests to AWS must be signed with an access key, which consists of an access key ID and secret access key. These two keys are commonly referred to as your security credentials. For more information about SigV4, see <u>Signature Version 4 signing process</u>.

Set up Amazon Managed Grafana for use with Amazon Managed Service for Prometheus

Amazon Managed Grafana is a fully managed service for open-source Grafana that simplifies connecting to open-source, third-party ISV, and AWS services for visualizing and analyzing your data sources at scale.

Amazon Managed Service for Prometheus supports using Amazon Managed Grafana to query metrics in a workspace. In the Amazon Managed Grafana console, you can add an Amazon Managed Service for Prometheus workspace as a data source by discovering your existing Amazon Managed Service for Prometheus accounts. Amazon Managed Grafana manages the configuration of the authentication credentials that are required to access Amazon Managed Service for Prometheus. For detailed instructions on creating a connection to Amazon Managed Service for Prometheus from Amazon Managed Grafana, see the instructions in <u>the Amazon Managed Grafana</u> <u>User Guide</u>.

You may also view your Amazon Managed Service for Prometheus alerts in Amazon Managed Grafana. For instructions to set up integration with alerts, see <u>Integrate alerts with Amazon</u> Managed Grafana or open source Grafana.

Connecting to Amazon Managed Grafana in a private VPC

Amazon Managed Service for Prometheus provides a service endpoint for Amazon Managed Grafana to connect to when querying metrics and alerts.

You can configure Amazon Managed Grafana to use a private VPC (for details on setting up a private VPC in Grafana, see <u>Connecting to Amazon VPC</u> in the *Amazon Managed Grafana User Guide*). Depending on the settings, this VPC may not have access to the Amazon Managed Service for Prometheus service endpoint.

To add Amazon Managed Service for Prometheus as a data source to an Amazon Managed Grafana workspace that is configured to use a specific private VPC, you must first connect your Amazon Managed Service for Prometheus to the same VPC by creating a VPC endpoint. For more information about creating a VPC endpoint, see <u>Create an interface VPC endpoint for Amazon</u> <u>Managed Service for Prometheus</u>.

Set up Grafana open source or Grafana Enterprise for use with Amazon Managed Service for Prometheus

You can use an instance of Grafana to query your metrics in Amazon Managed Service for Prometheus. This topic takes you through how to query metrics from Amazon Managed Service for Prometheus using a standalone instance of Grafana.

Prerequisites

Grafana instance – You must have a Grafana instance that is capable of authenticating with Amazon Managed Service for Prometheus.

Amazon Managed Service for Prometheus supports the use of Grafana version 7.3.5 and later to query metrics in a workspace. Versions 7.3.5 and later include support for AWS Signature Version 4 (SigV4) authentication.

To check your Grafana version, enter the following command, replacing *grafana_install_directory* with the path to your Grafana installation:

grafana_install_directory/bin/grafana-server -v

If you do not already have a standalone Grafana, or need a newer version, you can install a new instance. For instructions to set up a standalone Grafana, see Install Grafana in the Grafana

documentation. For information about getting started with Grafana, see <u>Getting started with</u> <u>Grafana</u> in the Grafana documentation.

AWS account – You must have an AWS account with the correct permissions to access your Amazon Managed Service for Prometheus metrics.

To set up Grafana to work with Amazon Managed Service for Prometheus, you must be logged on to an account that has the **AmazonPrometheusQueryAccess** policy or the aps:QueryMetrics, aps:GetMetricMetadata, aps:GetSeries, and aps:GetLabelspermissions. For more information, see <u>IAM permissions and policies</u>.

The next section describes setting up authentication from Grafana in more detail.

Step 1: Set up AWS SigV4

Amazon Managed Service for Prometheus works with AWS Identity and Access Management (IAM) to secure all calls to Prometheus APIs with IAM credentials. By default, the Prometheus data source in Grafana assumes that Prometheus requires no authentication. To enable Grafana to take advantage of Amazon Managed Service for Prometheus authentication and authorization capabilities, you will need to enable SigV4 authentication support in the Grafana data source. Follow the steps on this page when you are using a self-managed Grafana open-source or a Grafana enterprise server. If you are using Amazon Managed Grafana, SIGv4 authentication is fully automated. For more information about Amazon Managed Grafana, see <u>What is Amazon Managed Grafana</u>?

To enable SigV4 on Grafana, start Grafana with the AWS_SDK_LOAD_CONFIG and GF_AUTH_SIGV4_AUTH_ENABLED environment variables set to true. The GF_AUTH_SIGV4_AUTH_ENABLED environment variable overrides the default configuration for Grafana to enable SigV4 support. For more information, see <u>Configuration</u> in the Grafana documentation.

Linux

To enable SigV4 on a standalone Grafana server on Linux, enter the following commands.

export AWS_SDK_LOAD_CONFIG=true

export GF_AUTH_SIGV4_AUTH_ENABLED=true

cd grafana_install_directory

./bin/grafana-server

Windows

To enable SigV4 on a standalone Grafana on Windows using the Windows command prompt, enter the following commands.

set AWS_SDK_LOAD_CONFIG=true

set GF_AUTH_SIGV4_AUTH_ENABLED=true

cd grafana_install_directory

.\bin\grafana-server.exe

Step 2: Add the Prometheus data source in Grafana

The following steps explain how to set up the Prometheus data source in Grafana to query your Amazon Managed Service for Prometheus metrics.

To add the Prometheus data source in your Grafana server

- 1. Open the Grafana console.
- 2. Under **Configurations**, choose **Data sources**.
- 3. Choose Add data source.
- 4. Choose **Prometheus**.
- 5. For the HTTP URL, specify the **Endpoint query URL** displayed in the workspace details page in the Amazon Managed Service for Prometheus console.
- 6. In the HTTP URL that you just specified, remove the /api/v1/query string that is appended to the URL, because the Prometheus data source will automatically append it.

The correct URL should look similar to https://aps-workspaces.us-west-2.amazonaws.com/ workspaces/ws-1234a5b6-78cd-901e-2fgh-3i45j6k178l9.

- 7. Under Auth, select the toggle for SigV4 Auth to enable it.
- 8. You can either configure SigV4 authorization by specifying your long-term credentials directly in Grafana, or by using a default provider chain. Specifying your long-term credentials directly gets you started quicker, and the following steps give those instructions first. Once you are more familiar with using Grafana with Amazon Managed Service for Prometheus, we recommend that you use a default provider chain, because it provides better flexibility and security. For more information about setting up your default provider chain, see <u>Specifying</u> <u>Credentials</u>.
 - To use your long-term credentials directly, do the following:
 - a. Under SigV4 Auth Details, for Authentication Provider choose Access & secret key.
 - b. For Access Key ID, enter your AWS access key ID.
 - c. For Secret Access Key, enter your AWS secret access key.
 - d. Leave the Assume Role ARN and External ID fields blank.
 - e. For **Default Region**, choose the Region of your Amazon Managed Service for Prometheus workspace. This Region should match the Region contained in the URL that you listed in step 5.
 - f. Choose **Save & Test**.

You should see the following message: Data source is working

The following screenshot shows the Access key, Secret key SigV4 auth detail setting.

SigV4 Auth Details				
Authentication Provider	(Access & secret key		
Access Key ID		Configured	Ø	
Secret Access Key		Configured	Ø	
Assume Role ARN	3	arn:aws:iam:*		
External ID	3	External ID		
Default Region	3	us-west-2		

• To use a default provider chain instead (recommended for a production environment), do the following:

- a. Under SigV4 Auth Details, for Authentication Provider choose AWS SDK Default.
- b. Leave the Assume Role ARN and External ID fields blank.
- c. For **Default Region**, choose the Region of your Amazon Managed Service for Prometheus workspace. This Region should match the Region contained in the URL that you listed in step 5.
- d. Choose Save & Test.

You should see the following message: Data source is working

If you do not see that message, the next section provides troubleshooting tips for connecting.

The following screenshot shows the SDK default SigV4 auth detail setting.

SigV4 Auth Details	
Authentication Provider ③	AWS SDK Default
Assume Role ARN ③	arn:aws:iam:*
External ID ③	External ID
Default Region ③	us-west-2

- 9. Test a PromQL query against the new data source:
 - a. Choose Explore.
 - b. Run a sample PromQL query such as:

prometheus_tsdb_head_series

Step 3: (optional) Troubleshooting if Save & Test doesn't work

In the previous procedure, if you see an error when you choose Save & Test, check the following.

HTTP Error Not Found

Make sure that the workspace ID in the URL is correct.

HTTP Error Forbidden

This error means that the credentials are not valid. Check the following:

- Check that the Region specified in **Default Region** is correct.
- Check your credential for typos.
- Make sure that the credential that you are using has the **AmazonPrometheusQueryAccess** policy. For more information, see <u>IAM permissions and policies</u>.
- Make sure that the credential that you are using has access to this Amazon Managed Service for Prometheus workspace.

HTTP Error Bad Gateway

Look at the Grafana server log to troubleshoot this error. For more information, see Troubleshooting in the Grafana documentation.

If you see Error http: proxy error: NoCredentialProviders: no valid providers in chain, the default credential provider chain was not able to find a valid AWS credential to use. Make sure you have set up your credentials as documented in <u>Specifying Credentials</u>. If you want to use a shared configuration, make sure that the AWS_SDK_LOAD_CONFIG environment is set to true.

Query using Grafana running in an Amazon EKS cluster

Amazon Managed Service for Prometheus supports the use of Grafana version 7.3.5 and later to query metrics in a Amazon Managed Service for Prometheus workspace. Versions 7.3.5 and later include support for AWS Signature Version 4 (SigV4) authentication.

To set up Grafana to work with Amazon Managed Service for Prometheus, you must be logged on to an account that has the **AmazonPrometheusQueryAccess** policy or the aps:QueryMetrics, aps:GetMetricMetadata, aps:GetSeries, and aps:GetLabels permissions. For more information, see IAM permissions and policies.

Set up AWS SigV4

Grafana has added a new feature to support AWS Signature Version 4 (SigV4) authentication. For more information, see <u>Signature Version 4 signing process</u>. This feature is not enabled by default

on Grafana servers. The following instructions for enabling this feature assume that you are using Helm to deploy Grafana on a Kubernetes cluster.

To enable SigV4 on your Grafana 7.3.5 or later server

- 1. Create a new update file to override your Grafana configuration, and name it amp_query_override_values.yaml.
- 2. Enter the following content into the file, and save the file. Replace *account-id* with the AWS account ID where the Grafana server is running.

```
serviceAccount:
    name: "amp-iamproxy-query-service-account"
    annotations:
        eks.amazonaws.com/role-arn: "arn:aws:iam::account-id:role/amp-iamproxy-
query-role"
grafana.ini:
    auth:
        sigv4_auth_enabled: true
```

In that YAML file content, amp-iamproxy-query-role is the name of the role that you will create in the next section, <u>Set up IAM roles for service accounts</u>. You can replace this role with your own role name if you already have a role created for querying your workspace.

You will use this file later, in <u>Upgrade the Grafana server using Helm</u>.

Set up IAM roles for service accounts

If you are using a Grafana server in an Amazon EKS cluster, we recommend that you use IAM roles for service accounts, also known as service roles, for your access control. When you do this to associate an IAM role with a Kubernetes service account, the service account can then provide AWS permissions to the containers in any pod that uses that service account. For more information, see IAM roles for service accounts.

If you have not already set up these service roles for querying, follow the instructions at <u>Set up IAM</u> roles for service accounts for the querying of metrics to set up the roles.

You then need to add the Grafana service account in the conditions of the trust relationship.

To add the Grafana service account in the conditions of the trust relationship

1. From a terminal window, determine the namespace and the service account name for your Grafana server. For example, you could use the following command.

```
kubectl get serviceaccounts -n grafana_namespace
```

- 2. In the Amazon EKS console, open the IAM role for service accounts that is associated with the EKS cluster.
- 3. Choose Edit trust relationship.
- 4. Update the **Condition** to include the Grafana namespace and the Grafana service account name that you found in the output of the command in step 1. The following is an example.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Federated": "arn:aws:iam::account-id:oidc-provider/
oidc.eks.aws_region.amazonaws.com/id/openid"
      },
      "Action": "sts:AssumeRoleWithWebIdentity",
      "Condition": {
        "StringEquals": {
          "oidc.eks.region.amazonaws.com/id/openid:sub": [
            "system:serviceaccount:aws-amp:amp-iamproxy-query-service-account",
            "system:serviceaccount:grafana-namespace:grafana-service-account-name"
          1
        }
      }
    }
  ]
}
```

5. Choose **Update trust policy**.

Upgrade the Grafana server using Helm

This step upgrades the Grafana server to use the entries that you added to the amp_query_override_values.yaml file in the previous section.

Run the following commands. For more information about Helm charts for Grafana, see <u>Grafana</u> Community Kubernetes Helm Charts.

```
helm repo add grafana https://grafana.github.io/helm-charts
```

```
helm upgrade --install grafana grafana/grafana -n grafana_namespace -f ./
amp_query_override_values.yaml
```

Add the Prometheus data source in Grafana

The following steps explain how to set up the Prometheus data source in Grafana to query your Amazon Managed Service for Prometheus metrics.

To add the Prometheus data source in your Grafana server

- 1. Open the Grafana console.
- 2. Under Configurations, choose Data sources.
- 3. Choose Add data source.
- 4. Choose **Prometheus**.
- 5. For the HTTP URL, specify the **Endpoint query URL** displayed in the workspace details page in the Amazon Managed Service for Prometheus console.
- 6. In the HTTP URL that you just specified, remove the /api/v1/query string that is appended to the URL, because the Prometheus data source will automatically append it.
- 7. Under Auth, select the toggle for SigV4 Auth to enable it.

Leave the **Assume Role ARN** and **External ID** fields blank. Then for **Default Region**, select the Region where your Amazon Managed Service for Prometheus workspace is.

8. Choose **Save & Test**.

You should see the following message: Data source is working

- 9. Test a PromQL query against the new data source:
 - a. Choose Explore.
 - b. Run a sample PromQL query such as:

prometheus_tsdb_head_series

Query using Prometheus-compatible APIs

Although using a tool such as <u>Amazon Managed Grafana</u> is the easiest way to view and query your metrics, Amazon Managed Service for Prometheus also supports several Prometheuscompatible APIs that you can use to query your metrics. For more information about all the available Prometheus-compatible APIs, see Prometheus-compatible APIs.

The Prometheus-compatible APIs use the Prometheus query language, PromQL, to specify the data that you want to return. For details about PromQL and its syntax, see <u>Querying Prometheus</u> in the Prometheus documentation.

When you use these APIs to query your metrics, the requests must be signed with the AWS Signature Version 4 signing process. You can set up <u>AWS Signature Version 4</u> to simplify the signing process. For more information, see <u>aws-sigv4-proxy</u>.

Signing through AWS SigV4 proxy can be performed using awscurl. The following topic <u>Using</u> <u>awscurl to query Prometheus-compatible APIs</u> walks you through using awscurl to set up AWS SigV4.

Topics

• Use awscurl to query with Prometheus-compatible APIs

Use awscurl to query with Prometheus-compatible APIs

API requests for Amazon Managed Service for Prometheus must be signed with <u>SigV4</u>. You can use awscurl to simplify the querying process.

To install awscurl, you need to have Python 3 and pip package manager installed.

On a Linux based instance, the following command installs awscurl.

```
$ pip3 install awscurl
```

On a macOS machine, the following command installs awscurl.

```
$ brew install awscurl
```

The following example is a sample awscurl query. Replace the *Region*, *Workspace-id* and *QUERY* inputs with appropriate values for your use case:

Note

Your query string must be url encoded.

For a query like query=up, you could get results such as:

```
{
  "status": "success",
  "data": {
    "resultType": "vector",
    "result": [
      {
        "metric": {
          "___name___": "up",
          "instance": "localhost:9090",
          "job": "prometheus",
          "monitor": "monitor"
        },
        "value": [
          1652452637.636,
          "1"
        ]
      },
    ]
  }
}
```

In order for awscurl to sign the provided requests, you will need to pass the valid credentials in one of the following ways:

• Provide the access key ID and the Secret key for the IAM role. You can find the access key and the secret key for the role in the https://console.aws.amazon.com/iam/.

For example:

• Reference the configuration files stored in the .aws/credentials and /aws/config file. You can also choose to specify the name of the profile to be used. If unspecified, the default file will be used. For example:

• Use the instance profile associated with the EC2 instance.

Executing query requests using awscurl container

When installing a different version of **Python** and the associated dependencies is not feasible, a container can be used to package the awscurl application and its dependencies. The following example uses a **Docker** runtime to deploy awscurl, but any OCI compliant runtime and image will work.

```
$ docker pull okigan/awscurl
$ export AMP_QUERY_ENDPOINT=https://aps-workspaces.Region.amazonaws.com/
workspaces/Workspace_id/api/v1/query
```

\$ docker run --rm -it okigan/awscurl --access_key \$AWS_ACCESS_KEY_ID --secret_key
\$AWS_SECRET_ACCESS_KEY \ --region Region --service aps "\$AMP_QUERY_ENDPOINT?
query=QUERY"

Get statistics about your query usage for each query

Query pricing is based on the total number of query samples processed in a month from executed queries. You can get statistics about each query that you make to keep track of your samples processed. The query response for a query or a queryRange API can include the statistics data about query samples processed by including the query parameter stats=all in the request. A samples object is created in the stats object and the stats data is returned in the response.

The samples object consists of the following attributes:

Attribute	Description
totalQueryableSamples	Total number of query samples processed. This is the information to be used for billing.
totalQueryableSamp lesPerStep	The number of query samples processed per each step. This is structured as an array of arrays with the timestamp in epoch and the number of samples loaded on the specific step.

Sample requests and responses that include the stats information in the response are as follows:

Example for query:

GET

endpoint/api/v1/query?query=up&time=1652382537&stats=all

Response

```
{
    "status": "success",
    "data": {
        "resultType": "vector",
```

```
"result": [
            {
                 "metric": {
                     "___name___": "up",
                     "instance": "localhost:9090",
                     "job": "prometheus"
                 },
                 "value": [
                     1652382537,
                     "1"
                 ]
            }
        ],
        "stats": {
            "timings": {
                 "evalTotalTime": 0.00453349,
                 "resultSortTime": 0,
                 "queryPreparationTime": 0.000019363,
                 "innerEvalTime": 0.004508405,
                 "execQueueTime": 0.000008786,
                 "execTotalTime": 0.004554219
            },
            "samples": {
                 "totalQueryableSamples": 1,
                 "totalQueryableSamplesPerStep": [
                     Ε
                         1652382537,
                         1
                     ]
                 ]
            }
        }
    }
}
```

Example for queryRange:

GET

endpoint/api/v1/query_range?query=sum+%28rate+%28go_gc_duration_seconds_count%5B1m%5D
%29%29&start=1652382537&end=1652384705&step=1000&stats=all

Response

```
{
    "status": "success",
    "data": {
        "resultType": "matrix",
        "result": [
            {
                 "metric": {},
                 "values": [
                     Γ
                         1652383000,
                         "0"
                     ],
                     Ε
                         1652384000,
                         "0"
                     ]
                 ]
            }
        ],
        "stats": {
             "samples": {
                 "totalQueryableSamples": 8,
                 "totalQueryableSamplesPerStep": [
                     Ε
                         1652382000,
                         0
                     ],
                     Γ
                         1652383000,
                         4
                     ],
                     Ε
                         1652384000,
                         4
                     ]
                 ]
            }
        }
    }
}
```

Using rules to modify or monitor metrics as they are received

You can set up rules to act upon metrics as they are received by Amazon Managed Service for Prometheus. These rules can monitor the metrics or even create new, computed, metrics based on the metrics received.

Amazon Managed Service for Prometheus supports two types of *rules* that it evaluates at regular intervals:

- *Recording rules* allow you to precompute frequently needed or computationally expensive expressions and save their results as a new set of time series. Querying the precomputed result is often much faster than running the original expression every time it is needed.
- Alerting rules allow you to define alert conditions based on PromQL and a threshold. When the
 rule triggers the threshold, a notification is sent to <u>alert manager</u>, which can be configured to
 managed the rules, or forward them to notification downstream to receivers such as Amazon
 Simple Notification Service.

To use rules in Amazon Managed Service for Prometheus, you create one or more YAML rules files that define the rules. An Amazon Managed Service for Prometheus rules file has the same format as a rules file in standalone Prometheus. For more information, see <u>Defining Recording rules</u> and <u>Alerting rules</u> in the Prometheus documentation.

You can have multiple rules files in a workspace. Each separate rules file is contained within a separate *namespace*. Having multiple rules files lets you import existing Prometheus rules files to a workspace without having to change or combine them. Different rule group namespaces can also have different tags.

Rule sequencing

Within a rules file, rules are contained within *rules groups*. Rules within a single rules group in a rules file are always evaluated in order from top to bottom. Therefore, in recording rules, the result of one recording rule can be used in the computation of a later recording rule or in an alerting rule in the same rule group. However, because you can't specify the order in which to run separate rules files, you can't use the results from one recording rule to compute a rule in a different rule group or a different rules file.

Topics

- Understanding IAM permissions needed for using rules
- Create a rules file
- Upload a rules configuration file to Amazon Managed Service for Prometheus
- Edit or replace a rules configuration file
- Troubleshooting Ruler

Understanding IAM permissions needed for using rules

You must give users permissions to use rules in Amazon Managed Service for Prometheus. Create an AWS Identity and Access Management (IAM) policy with the following permissions, and assign the policy to your users, groups, or roles.

1 Note

For more information about IAM, see <u>Identity and Access Management for Amazon</u> Managed Service for Prometheus.

Policy to give access to use rules

The following policy gives access to use rules for all resources in your account.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
               "aps: CreateRuleGroupsNamespace",
               "aps: ListRuleGroupsNamespaces",
               "aps: DescribeRuleGroupsNamespace",
               "aps: PutRuleGroupsNamespace",
               "aps: DeleteRuleGroupsNamespace",
                "aps: DeleteRuleGroupsNamespace",
               "aps: DeleteRuleGroupsNamespace",
                "aps: DeleteRuleGroupsNamespace",
                "aps: DeleteRuleGroupsNamespace",
               "aps: DeleteRuleGroupsNamespace",
               "aps: DeleteRuleGroupsNamespace",
               "aps: DeleteRuleGroupsNamespace",
               "aps: DeleteRuleGroupsNamespace",
```

}

Policy to give access to only one namespace

You can also create policy that gives access to only specific policies. The following sample policy gives access only to the RuleGroupNamespace specified. To use this policy, replace <a count>, <region>, <workspace-id>, and <namespace-name> with appropriate values for your account.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "aps:ListRules",
                "aps:ListTagsForResource",
                "aps:GetLabels",
                "aps:CreateRuleGroupsNamespace",
                "aps:ListRuleGroupsNamespaces",
                "aps:DescribeRuleGroupsNamespace",
                "aps:PutRuleGroupsNamespace",
                "aps:DeleteRuleGroupsNamespace"
            ],
            "Resource": [
                "arn:aws:aps:*:<account>:workspace/*",
                "arn:aws:aps:<region>:<account>:rulegroupnamespace/<workspace-
id>/<namespace-name>"
            ]
        }
    ]
}
```

Create a rules file

To use rules in Amazon Managed Service for Prometheus, you create a rules file that defines the rules. An Amazon Managed Service for Prometheus rules file is a YAML text file that has the same format as a rules file in standalone Prometheus. For more information, see <u>Defining Recording</u> rules and <u>Alerting rules</u> in the *Prometheus* documentation.

The following is a basic example of a rules file:

```
groups:
        - name: cpu_metrics
        rules:
        - record: avg_cpu_usage
        expr: avg(rate(node_cpu_seconds_total[5m])) by (instance)
        - alert: HighAverageCPU
        expr: avg_cpu_usage > 0.8
        for: 10m
        keep_firing_for: 20m
        labels:
            severity: critical
        annotations:
            summary: "Average CPU usage across cluster is too high"
```

This simple example creates a new metric using a recording rule, called avg_cpu_usage and then uses that in an alert. The following describes some of the properties used. For more information about alerting rules and other properties you can include, see <u>Alerting rules</u> in the *Prometheus* documentation.

- record: avg_cpu_usage This recording rule creates a new metric called avg_cpu_usage.
- expr: avg(rate(node_cpu_seconds_total[5m])) by (instance) This expression for the recording rule calculates the average rate of CPU usage over the last 5 minutes for each node, grouping by the instance label.
- alert: HighAverageCPU This alert rule creates a new alert called HighAverageCPU
- expr: avg_cpu_usage > 0.8 This expression tells the alert to look for samples where the average CPU usage goes over 80%.
- for: 10m The alert will fire when the expression is met for 10 minutes. In this case, the samples are an average over 5 minutes, so the alert will fire when it receives at least 2 samples that are over the threshold.
- keep_firing_for: 20m This alert will continue to fire until the samples are below the threshold for at least 20 minutes. This can be useful to avoid the alert going up and down repeatedly in succession.

For more alerting rule examples, see <u>Alerting rule examples</u>.

🚯 Note

You can create a rules definition file locally and then upload it to Amazon Managed Service for Prometheus, or you can create, edit and upload the definition directly within the Amazon Managed Service for Prometheus console. Either way, the same formatting rules apply. To learn more about uploading and editing your file, see <u>Upload a rules</u> configuration file to Amazon Managed Service for Prometheus.

Upload a rules configuration file to Amazon Managed Service for Prometheus

Once you know what rules you want in your rules configuration file, you can either create and edit it within the console, or you can upload a file with the console or AWS CLI.

🚯 Note

If you are running an Amazon EKS cluster, you can also upload a rule configuration file using <u>AWS Controllers for Kubernetes</u>.

To use the Amazon Managed Service for Prometheus console to edit or replace your rules configuration and create the namespace

- 1. Open the Amazon Managed Service for Prometheus console at https://console.aws.amazon.com/prometheus/.
- 2. In the upper left corner of the page, choose the menu icon, and then choose All workspaces.
- 3. Choose the workspace ID of the workspace, and then choose the **Rules management** tab.
- 4. Choose Add namespace.
- 5. Choose **Choose file**, and select the rules definition file.

Alternately, you can create and edit a rules definition file directly in the Amazon Managed Service for Prometheus console by selecting **Define configuration**. This will create a sample default definition file that you edit before uploading.

6. (Optional) To add tags to the namespace, choose Add new tag.

Then, for Key, enter a name for the tag. You can add an optional value for the tag in Value.

To add another tag, choose Add new tag.

7. Choose **Continue**. Amazon Managed Service for Prometheus creates a new namespace with the same name as the rules file that you selected.

To use the AWS CLI to upload an alert manager configuration to a workspace in a new namespace

1. Base64 encode the contents of your alert manager file. On Linux, you can use the following command:

base64 input-file output-file

On macOS, you can use the following command:

openssl base64 input-file output-file

2. Enter one of the following commands to create the namespace and upload the file.

On AWS CLI version 2, enter:

```
aws amp create-rule-groups-namespace --data file://path_to_base_64_output_file --
name namespace-name --workspace-id my-workspace-id --region region
```

On AWS CLI version 1, enter:

```
aws amp create-rule-groups-namespace --data fileb://path_to_base_64_output_file --
name namespace-name --workspace-id my-workspace-id --region region
```

3. It takes a few seconds for your alert manager configuration to become active. To check the status, enter the following command:

```
aws amp describe-rule-groups-namespace --workspace-id workspace_id --
name namespace-name --region region
```

If the status is ACTIVE, your rules file has taken effect.

If you want to change the rules in a rule file that you have already uploaded to Amazon Managed Service for Prometheus, you can either upload a new rules file to replace the existing configuration, or you can edit the current configuration directly in the console. Optionally, you can download the current file, edit it in a text editor, then upload the new version.

To use the Amazon Managed Service for Prometheus console to edit your rules configuration

- 1. Open the Amazon Managed Service for Prometheus console at https://console.aws.amazon.com/prometheus/.
- 2. In the upper left corner of the page, choose the menu icon, and then choose All workspaces.
- 3. Choose the workspace ID of the workspace, and then choose the **Rules management** tab.
- 4. Select the name of the rules configuration file that you want to edit.
- (Optional) If you want to download the current rules configuration file, choose **Download** or **Copy**.
- 6. Choose **Modify** to edit the configuration directly within the console. Choose **Save** when complete.

Alternately, you can choose **Replace configuration** to upload a new configuration file. If so, select the new rules definition file, and choose **Continue** to upload it.

To use the AWS CLI to edit a rules configuration file

1. Base64 encode the contents of your rules file. On Linux, you can use the following command:

base64 input-file output-file

On macOS, you can use the following command:

openssl base64 input-file output-file

2. Enter one of the following commands to upload the new file.

On AWS CLI version 2, enter:

```
aws amp put-rule-groups-namespace --data file://path_to_base_64_output_file --
name namespace-name --workspace-id my-workspace-id --region region
```

On AWS CLI version 1, enter:

```
aws amp put-rule-groups-namespace --data fileb://path_to_base_64_output_file --
name namespace-name --workspace-id my-workspace-id --region region
```

3. It takes a few seconds for your rules file to become active. To check the status, enter the following command:

```
aws amp describe-rule-groups-namespace --workspace-id workspace_id --
name namespace-name --region region
```

If the status is ACTIVE, your rules file has taken effect. Until then, the previous version of this rules file is still active.

Troubleshooting Ruler

Using <u>Monitor Amazon Managed Service for Prometheus events with CloudWatch Logs</u>, you can troubleshoot Alert Manager and Ruler related issues. This section contains ruler related troubleshooting topics.

When the log contains the following ruler failure error

```
{
    "workspaceId": "ws-12345c67-89c0-4d12-345b-f14db70f7a99",
    "message": {
        "log": "Evaluating rule failed, name=failure,
    group=canary_long_running_vl_namespace, namespace=canary_long_running_vl_namespace,
    err=found duplicate series for the match group {dimension1=\\\"1\\\"} on the right
    hand-side of the operation: [{___name__=\\\"fake_metric2\\\", dimension1=\\\"1\\\",
    dimension2=\\\"b\\\"}, {___name__=\\\"fake_metric2\\\", dimension1=\\\"1\\\",
    dimension2=\\\"a\\\"}];many-to-many matching not allowed: matching labels must be
    unique on one side",
        "level": "ERROR",
        "name": "failure",
        "group": "canary_long_running_vl_namespace",
    }
}
```

```
"namespace": "canary_long_running_vl_namespace"
},
"component": "ruler"
}
```

This means that some error occurred while executing the rule.

Action to take

Use the error message to troubleshoot the rule execution.

Managing and forwarding alerts in Amazon Managed Service for Prometheus with alert manager

When the <u>alerting rules</u> that Amazon Managed Service for Prometheus runs are firing, alert manager handles the alerts that are sent. It deduplicates, groups, and routes the alerts to downstream receivers. Amazon Managed Service for Prometheus supports only Amazon Simple Notification Service as a receiver, and can route messages to Amazon SNS topics in the same account. You can also use alert manager to silence and inhibit alerts.

Alert manager provides similar functionality to Alertmanager in Prometheus.

You can use alert manager's configuration file for the following:

 Grouping – Grouping collects similar alerts into a single notification. This is especially useful during larger outages when many systems fail at once and hundreds of alerts might fire simultaneously. For example, suppose that a network failure causes many of your nodes to fail at the same time. If these types of alerts are grouped, alert manager sends you a single notification.

Alert grouping and the timing for the grouped notifications are configured by a routing tree in the alert manager configuration file. For more information, see .

- Inhibition Inhibition suppresses notifications for certain alerts if certain other alerts are already firing. For example, if an alert is firing about a cluster being unreachable, you can configure alert manager to mute all other alerts concerning this cluster. This prevents notifications for hundreds or thousands of firing alerts that are unrelated to the actual issue. For more information about how to write inhibition rules, see <inhibit_rule>.
- **Silences** Silences mute alerts for a specified time, such as during a maintenance window. Incoming alerts are checked for whether they match all the equality or regular expression matchers of an active silence. If they do, no notifications are sent for that alert.

To create a silence, you use the PutAlertManagerSilences API. For more information, see PutAlertManagerSilences.

Prometheus templating

Standalone Prometheus supports templating, using separate template files. Templates can use conditionals and format data, among other things.

In Amazon Managed Service for Prometheus, you put your templating in the same alert manager configuration file as your alert manager configuration.

Topics

- Understanding IAM permissions needed for working with alert manager
- <u>Create an alert manager configuration in Amazon Managed Service for Prometheus to manage</u> and route alerts
- Forward alerts to an alert receiver with alert manager in Amazon Managed Service for Prometheus
- Upload your alert manager configuration file to Amazon Managed Service for Prometheus
- Integrate alerts with Amazon Managed Grafana or open source Grafana
- <u>Troubleshoot alert manager with CloudWatch Logs</u>

Understanding IAM permissions needed for working with alert

manager

You must give users permissions to use alert manager in Amazon Managed Service for Prometheus. Create an AWS Identity and Access Management (IAM) policy with the following permissions, and assign the policy to your users, groups, or roles.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "aps: CreateAlertManagerDefinition",
                "aps: DescribeAlertManagerSilence",
                "aps: DescribeAlertManagerDefinition",
                "aps: PutAlertManagerDefinition",
                "aps: DeleteAlertManagerDefinition",
                "aps: ListAlerts",
                "aps: ListRules",
                "aps: ListAlertManagerReceivers",
                "aps: ListAlertManagerSilences",
                "aps: ListAlertManagerAlerts",
                "aps: ListAlertManagerAlertGroups",
                "aps: GetAlertManagerStatus",
```

}

```
"aps: GetAlertManagerSilence",
    "aps: PutAlertManagerSilences",
    "aps: DeleteAlertManagerSilence",
    "aps: CreateAlertManagerAlerts"
],
    "Resource": "*"
}
```

Create an alert manager configuration in Amazon Managed Service for Prometheus to manage and route alerts

To use alert manager and templating in Amazon Managed Service for Prometheus, you create an alert manager configuration YAML file. An Amazon Managed Service for Prometheus alert manager file has two main sections:

- template_files: contains the templates used for messages sent by receivers. For more information, see Template Reference and Template Examples in the Prometheus documentation.
- alertmanager_config: contains the alert manager configuration. This uses the same structure as an alert manager config file in standalone Prometheus. For more information, see <u>Configuration</u> in the Alertmanager documentation.

Note

The repeat_interval configuration described in the Prometheus documentation above has an additional limitation in Amazon Managed Service for Prometheus. The maximum allowed value is five days. If you set it higher than five days, it will be treated as five days and notifications will be sent again after the five day period has passed.

🚯 Note

You can also edit the configuration file directly in the Amazon Managed Service for Prometheus console, but it must still follow the format specified here. For more information on uploading or editing a configuration file, see <u>Upload your alert manager</u> configuration file to Amazon Managed Service for Prometheus. In Amazon Managed Service for Prometheus, your alert manager configuration file must have all your alert manager configuration content inside of an alertmanager_config key at the root of the YAML file.

The following is a basic example alert manager configuration file:

```
alertmanager_config: |
route:
receiver: 'default'
receivers:
    name: 'default'
    sns_configs:
    topic_arn: arn:aws:sns:us-east-2:123456789012:My-Topic
    sigv4:
    region: us-east-2
    attributes:
        key: key1
        value: value1
```

The only receiver currently supported is Amazon Simple Notification Service (Amazon SNS). If you have other types of receivers listed in the configuration, it will be rejected.

Here is another sample alert manager configuration file that uses both the template_files block and the alertmanager_config block.

```
template_files:
  default_template: |
    {{ define "sns.default.subject" }}[{{ .Status | toUpper }}{{ if eq .Status
 "firing" }}:{{ .Alerts.Firing | len }}{{ end }}]{{ end }}
    {{ define "__alertmanager" }}AlertManager{{ end }}
    {{ define "__alertmanagerURL" }}{{ .ExternalURL }}/#/alerts?receiver={{ .Receiver |
 urlquery }}{{ end }}
alertmanager_config: |
  global:
  templates:
    - 'default_template'
  route:
    receiver: default
  receivers:
    - name: 'default'
      sns_configs:
      - topic_arn: arn:aws:sns:us-east-2:accountid:My-Topic
```

```
sigv4:
    region: us-east-2
attributes:
    key: severity
    value: SEV2
```

Default Amazon SNS template block

The default Amazon SNS configuration uses the following template unless you explicitly override it.

Forward alerts to an alert receiver with alert manager in Amazon Managed Service for Prometheus

When an alert is raised by an alert rule, it is sent to alert manager. Alert manager performs functions such as deduplicating alerts, inhibiting alerts during maintenance, or grouping them as needed. It then forwards the alert as a message to an *alert receiver*. You can set up an alert receiver that can notify operators, have automated responses, or respond to the alerts in other ways.

The only alert receiver supported in Amazon Managed Service for Prometheus is Amazon Simple Notification Service (Amazon SNS). For more information, see <u>What is Amazon SNS?</u>. Amazon SNS can be used to respond to alerts in a wide variety of ways, including forwarding to other systems, such as email, SMS, or HTTP endpoints.

The following topics describe the tasks associated with creating and configuring your Amazon SNS alert receiver.

Topics

- <u>Creating a new Amazon SNS topic for use as an alert receiver in Amazon Managed Service for</u>
 Prometheus
- <u>Giving Amazon Managed Service for Prometheus permission to send alert messages to your</u> <u>Amazon SNS topic</u>
- Configure alert manager to send messages to your Amazon SNS topic
- <u>Configure alert manager to send messages to Amazon SNS as JSON</u>
- Configure Amazon SNS to send messages for alerts to other destinations
- Understanding Amazon SNS message validation rules

Creating a new Amazon SNS topic for use as an alert receiver in Amazon Managed Service for Prometheus

You can use an existing Amazon SNS topic as an alert receiver for Amazon Managed Service for Prometheus, or you can create a new one. We recommend that you use a topic of the **Standard** type, so that you can forward alerts from the topic to email, SMS, or HTTP.

To create a new Amazon SNS topic to use as your alert manager receiver, follow the steps in <u>Step</u> <u>1: Create a topic</u>. Be sure to choose **Standard** for the topic type.

If you want to receive emails every time a message is sent to that Amazon SNS topic, follow the steps in <u>Step 2: Create a subscription to the topic</u>.

Whether you use a new or existing Amazon SNS topic, you will need the Amazon Resource Name (ARN) of your Amazon SNS topic to complete the following tasks.

Giving Amazon Managed Service for Prometheus permission to send alert messages to your Amazon SNS topic

You must give Amazon Managed Service for Prometheus permission to send messages to your Amazon SNS topic. The following policy statement will give that permission. It includes a Condition statement to help prevent a security problem known as the *confused deputy* problem. The Condition statement restricts access to the Amazon SNS topic to allow only operations coming from this specific account and Amazon Managed Service for Prometheus workspace. For more information about the confused deputy problem, see <u>Cross-service confused deputy</u> <u>prevention</u>.

To give Amazon Managed Service for Prometheus permission to send messages to your Amazon SNS topic

- 1. Open the Amazon SNS console at https://console.aws.amazon.com/sns/v3/home.
- 2. In the navigation pane, choose **Topics**.
- 3. Choose the name of the topic that you are using with Amazon Managed Service for Prometheus.
- 4. Choose Edit.
- 5. Choose **Access policy** and add the following policy statement to the existing policy.

```
{
    "Sid": "Allow_Publish_Alarms",
    "Effect": "Allow",
    "Principal": {
        "Service": "aps.amazonaws.com"
    },
    "Action": [
        "sns:Publish",
        "sns:GetTopicAttributes"
    ],
    "Condition": {
        "ArnEquals": {
            "aws:SourceArn": "workspace_ARN"
        },
        "StringEquals": {
            "AWS:SourceAccount": "account_id"
        }
    },
    "Resource": "arn:aws:sns:region:account_id:topic_name"
}
```

[Optional] If your Amazon SNS topic is service side encryption (SSE) enabled, you need to allow Amazon Managed Service for Prometheus to send messages to this encrypted topic by adding the kms:GenerateDataKey* and kms:Decrypt permissions to the AWS KMS key policy of the key used to encrypt the topic.

For example, you could add the following to the policy:

```
{
    "Statement": [{
```

```
"Effect": "Allow",
"Principal": {
    "Service": "aps.amazonaws.com"
    },
    "Action": [
        "kms:GenerateDataKey*",
        "kms:Decrypt"
    ],
    "Resource": "*"
  }]
}
```

For more information, see AWS KMS Permissions for SNS Topic.

6. Choose Save changes.

Note

By default, Amazon SNS creates the access policy with condition on AWS: SourceOwner. For more information, see <u>SNS Access Policy</u>.

🚺 Note

IAM follows the <u>Most-restrictive policy first</u> rule. In your SNS topic, if there is a policy block that is more restrictive than the documented Amazon SNS policy block, the permission for the topic policy is not granted. To evaluate your policy and find out what's been granted, see <u>Policy evaluation logic</u>.

Cross-service confused deputy prevention

The confused deputy problem is a security issue where an entity that doesn't have permission to perform an action can coerce a more-privileged entity to perform the action. In AWS, cross-service impersonation can result in the confused deputy problem. Cross-service impersonation can occur when one service (the *calling service*) calls another service (the *called service*). The calling service can be manipulated to use its permissions to act on another customer's resources in a way it should not otherwise have permission to access. To prevent this, AWS provides tools that help you protect

We recommend using the <u>aws:SourceArn</u> and <u>aws:SourceAccount</u> global condition context keys in resource policies to limit the permissions that Amazon Managed Service for Prometheus gives to Amazon SNS to the resource. If you use both global condition context keys, the aws:SourceAccount value and the account in the aws:SourceArn value must use the same account ID when used in the same policy statement.

The value of aws:SourceArn must be the ARN of the Amazon Managed Service for Prometheus workspace.

The most effective way to protect against the confused deputy problem is to use the aws:SourceArn global condition context key with the full ARN of the resource. If you don't know the full ARN of the resource or if you are specifying multiple resources, use the aws:SourceArn global context condition key with wildcards (*) for the unknown portions of the ARN. For example, arn:aws:servicename::123456789012:*.

The policy shown in <u>Giving Amazon Managed Service for Prometheus permission to send</u> <u>alert messages to your Amazon SNS topic</u> shows how you can use the aws:SourceArn and aws:SourceAccount global condition context keys in Amazon Managed Service for Prometheus to prevent the confused deputy problem.

Configure alert manager to send messages to your Amazon SNS topic

After you have a (new or existing) **Standard** type Amazon SNS topic, you can add it to your alert manager configuration as an alert receiver. Alert manager can forward your alerts to a configured alert receiver. To complete this, you must know the Amazon Resource Name (ARN) of your Amazon SNS topic.

For more information about Amazon SNS receiver configuration, see <u><sns_configs></u> in the Prometheus configuration documentation.

Unsupported properties

Amazon Managed Service for Prometheus supports Amazon SNS as the alert receiver. However, because of service constraints, not all of the properties of the Amazon SNS receiver are supported. The following properties are not allowed in an Amazon Managed Service for Prometheus alert manager configuration file:

- api_url: Amazon Managed Service for Prometheus sets the api_url for you, so this property is not allowed.
- Http_config This property allows you to set external proxies. Amazon Managed Service for Prometheus does not currently support this feature.

Additionally, SigV4 settings are required to have a Region property. Without the Region property, Amazon Managed Service for Prometheus doesn't have enough information to make the authorization request.

To configure alert manager with your Amazon SNS topic as the receiver

- 1. If you are using an existing alert manager configuration file, open it in a text editor.
- 2. If there are current receivers other than Amazon SNS in the receivers block, remove them. You can configure multiple Amazon SNS topics to be receivers by putting them in separate sns_config blocks within the receivers block.
- 3. Add the following YAML block within the receivers section.



If a subject is not specified, by default, a subject would be generated with the default template with the label name and values, which may result in a value that is too long for SNS. To change the template that is applied to the subject, refer to <u>Configure alert manager to send messages to</u> <u>Amazon SNS as JSON</u> in this guide.

Now you must upload your alert manager configuration file to Amazon Managed Service for Prometheus. For more information, see <u>Upload your alert manager configuration file to Amazon</u> <u>Managed Service for Prometheus</u>.

Configure alert manager to send messages to Amazon SNS as JSON

By default, Amazon Managed Service for Prometheus alert manager outputs messages in a plain text list format. This can be more difficult for other services to parse. You can configure alert manager to send alerts in JSON format instead. JSON can make it simpler to process the messages downstream from Amazon SNS in AWS Lambda or in webhook-receiving endpoints. Instead of using the default template, you can define a custom template to output the message contents in JSON, making it easier to parse in downstream functions.

To output messages from alert manager to Amazon SNS in JSON format, update your alert manager configuration to contain the following code inside your template_files root section:

```
default_template: |
  {{ define "sns.default.message" }}{{ "{" }}"receiver": "{{ .Receiver }}","status":
 "{{ .Status }}","alerts": [{{ range $alertIndex, $alerts := .Alerts }}{{ if
$alertIndex }}, {{ end }}{{ "{" }}"status": "{{ $alerts.Status }}"{{ if
gt (len $alerts.Labels.SortedPairs) 0 -}},"labels": {{ "{" }}{{ range
$index, $label := $alerts.Labels.SortedPairs }}{{ if $index }},
{{ end }}"{{ $label.Name }}": "{{ $label.Value }}"{{ end }}
{{ "}" }}{{- end }}{{ if gt (len $alerts.Annotations.SortedPairs )
0 -}},"annotations": {{ "{" }}{{ range $index, $annotations :=
$alerts.Annotations.SortedPairs }}{{ if $index }}, {{ end }}"{{ $annotations.Name }}":
"{{ $annotations.Value }}"{{ end }}{{ "}" }}{{- end }},"startsAt":
"{{ $alerts.StartsAt }}","endsAt": "{{ $alerts.EndsAt }}","generatorURL":
 "{{ $alerts.GeneratorURL }}","fingerprint": "{{ $alerts.Fingerprint }}"{{ "}" }}
{{ end }}]{{ if gt (len .GroupLabels) 0 -}},"groupLabels": {{ "{" }}{{ range
$index, $groupLabels := .GroupLabels.SortedPairs }}{{ if $index }},
{{ end }}"{{ $groupLabels.Name }}": "{{ $groupLabels.Value }}"{{ end }}
{{ "}" }}{{- end }}{{ if gt (len .CommonLabels) 0 -}},"commonLabels": {{ "{" }}
{{ range $index, $commonLabels := .CommonLabels.SortedPairs }}{{ if $index }},
{{ end }}"{{ $commonLabels.Name }}": "{{ $commonLabels.Value }}"{{ end }}{{ "}" }}{{-
end }}{{ if gt (len .CommonAnnotations) 0 -}},"commonAnnotations": {{ "{" }}{{ range
$index, $commonAnnotations := .CommonAnnotations.SortedPairs }}{{ if $index }},
{{ end }}"{{ $commonAnnotations.Name }}": "{{ $commonAnnotations.Value }}"{{ end }}
{{ "}" }}{{ - end }}{{ "}" }}{{ end }}
  {{ define "sns.default.subject" }}[{{ .Status | toUpper }}{{ if eq .Status
 "firing" }}:{{ .Alerts.Firing | len }}{{ end }}]{{ end }}
```

🚯 Note

This template creates JSON from alphanumeric data. If your data has special characters, encode them before using this template.

To make sure that this template is used in outgoing notifications, reference it in your alertmanager_config block as follows:

```
alertmanager_config: |
  global:
  templates:
    - 'default_template'
```

Note

This template is for the entire message body as JSON. This template overwrites the entire message body. You cannot override the message body if you wish to use this specific template. Any overrides that are manually done will take precedence over the template.

For more information about:

- The alert manager configuration file, see <u>Create an alert manager configuration in Amazon</u> Managed Service for Prometheus to manage and route alerts.
- Uploading your configuration file, see <u>Upload your alert manager configuration file to Amazon</u> <u>Managed Service for Prometheus</u>.

Configure Amazon SNS to send messages for alerts to other destinations

Amazon Managed Service for Prometheus can only send alert messages to Amazon Simple Notification Service (Amazon SNS). To send those messages to other destinations, such as email, webhook, Slack, or OpsGenie, you must configure Amazon SNS to forward the messages on to those endpoints.

The following sections describing configuring Amazon SNS to forward alerts to other destinations.

Topics

- Email
- Webhook
- <u>Slack</u>
- OpsGenie

Email

To configure an Amazon SNS topic to output messages to email, create a subscription. In the Amazon SNS console, choose the **Subscriptions** tab to open the **Subscriptions** list page. Choose **Create Subscription** and select **Email**. Amazon SNS sends a confirmation email to the listed email address. After you accept the confirmation, you are able to receive Amazon SNS notifications as emails from the topic you subscribed to. For more information, see <u>Subscribing to an Amazon SNS</u> topic.

Webhook

To configure an Amazon SNS topic to output messages to a webhook endpoint, create a subscription. In the Amazon SNS console, choose the **Subscriptions** tab to open the **Subscriptions** list page. Choose **Create Subscription** and select **HTTP/HTTPS**. After you create the subscription, you must follow the confirmation steps to activate it. When it is active, your HTTP endpoint should receive the Amazon SNS notifications. For more information, see <u>Subscribing to an Amazon</u> <u>SNS topic</u>. For more information about using Slack webhooks to publish messages to various destinations, see <u>How do I use webhooks to publish Amazon SNS messages to Amazon Chime</u>, Slack, or Microsoft Teams?

Slack

To configure an Amazon SNS topic to output messages to Slack, you have two options. You can either integrate with Slack's email-to-channel integration, which allows Slack to accept email messages and forward them to a Slack channel, or you can use a Lambda function to rewrite the Amazon SNS notification to Slack. For more information about forwarding emails to slack channels, see <u>Confirming AWS SNS Topic Subscription for Slack Webhook</u>. For more information about constructing a Lambda function to convert Amazon SNS messages to Slack, see <u>How to integrate Amazon Managed Service for Prometheus with Slack</u>.

OpsGenie

For information about how to configure an Amazon SNS topic to output messages to OpsGenie, see Integrate Opsgenie with Incoming Amazon SNS.

Understanding Amazon SNS message validation rules

Amazon Simple Notification Service (Amazon SNS) requires messages to meet certain standards. Messages that don't meed these standards will be modified when they are received. The alert messages will be validated, truncated, or modified, if necessary, by the Amazon SNS receiver based on the following rules:

- Message contains non-utf characters.
 - Message will be replaced by "Error not a valid UTF-8 encoded string."
 - One message attribute will be added with the key of "truncated" and the value of "true"
 - One message attribute will be added with the key of "modified" and the value of "Message: Error - not a valid UTF-8 encoded string."
- Message is empty.
 - Message will be replaced by "Error Message should not be empty."
 - One message attribute will be added with the key of "modified" and the value of "Message: Error - Message should not be empty."
- Message has been truncated.
 - Message will have the truncated content.
 - One message attribute will be added with the key of "truncated" and the value of "true"
 - One message attribute will be added with the key of "modified" and the value of "Message: Error - Message has been truncated from *X* KB, because it exceeds the 256 KB size limit."
- Subject is not ASCII.
 - Subject will be replaced by "Error contains non printable ASCII characters."
 - One message attribute will be added with the key of "modified" and the value of "Subject: Error - contains non-printable ASCII characters."
- Subject has been truncated.
 - Subject will have the truncated content.

- One message attribute will be added with the key of "modified" and the value of "Subject: Error - Subject has been truncated from X characters, because it exceeds the 100 character size limit."
- Message attribute has invalid key/value.
 - Invalid message attribute will be removed.
 - One message attribue will be added with the key of "modified" and the value of "MessageAttribute: Error - X of the message attributes have been removed becasue of invalid MessageAttributeKey or MessageAttributeValue."
- Message attribute has been truncated.
 - Extra message attributes will be removed.
 - One message attribute will be added with the key of "modified" and the value of "MessageAttribute: Error - X of the message attributes have been removed, because it exceeds the 256KB size limit.

Upload your alert manager configuration file to Amazon Managed Service for Prometheus

Once you know what you want in your Alert manager configuration file, you can create and edit it within the console, or you can upload an existing file with the Amazon Managed Service for Prometheus console or AWS CLI.

🚯 Note

If you are running an Amazon EKS cluster, you can also upload an Alert manager configuration file using AWS Controllers for Kubernetes.

To use the Amazon Managed Service for Prometheus console to edit or replace your alert manager configuration

- 1. Open the Amazon Managed Service for Prometheus console at <u>https://</u> console.aws.amazon.com/prometheus/.
- 2. In the upper left corner of the page, choose the menu icon, and then choose **All workspaces**.
- 3. Choose the workspace ID of the workspace, and then choose the **Alert manager** tab.

4. If the workspace doesn't already have an alert manager definition, choose **Add definition**.

Note

If the workspace has an alert manager definition that you want to replace, choose **Modify** instead.

5. Choose **Choose file**, select the alert manager definition file, and choose **Continue**.

🚯 Note

Alternately, you can create a new file and edit it directly in the console, by choosing the **Create definition** option. This will create a sample default configuration that you edit before uploading.

To use the AWS CLI to upload an alert manager configuration to a workspace for the first time

1. Base64 encode the contents of your alert manager file. On Linux, you can use the following command:

base64 input-file output-file

On macOS, you can use the following command:

openssl base64 input-file output-file

2. To upload the file, enter one of the following commands.

On AWS CLI version 2, enter:

```
aws amp create-alert-manager-definition --data file://path_to_base_64_output_file
    --workspace-id my-workspace-id --region region
```

On AWS CLI version 1, enter:

```
aws amp create-alert-manager-definition --data fileb://path_to_base_64_output_file
    --workspace-id my-workspace-id --region region
```

3. It takes a few seconds for your alert manager configuration to become active. To check the status, enter the following command:

```
aws amp describe-alert-manager-definition --workspace-id workspace_id --
region region
```

If the status is ACTIVE, your new alert manager definition has taken effect.

To use the AWS CLI to replace a workspace's alert manager configuration with a new one

1. Base64 encode the contents of your alert manager file. On Linux, you can use the following command:

base64 input-file output-file

On macOS, you can use the following command:

openssl base64 input-file output-file

2. To upload the file, enter one of the following commands.

On AWS CLI version 2, enter:

```
aws amp put-alert-manager-definition --data file://path_to_base_64_output_file --
workspace-id my-workspace-id --region region
```

On AWS CLI version 1, enter:

```
aws amp put-alert-manager-definition --data fileb://path_to_base_64_output_file --
workspace-id my-workspace-id --region region
```

3. It takes a few seconds for your new alert manager configuration to become active. To check the status, enter the following command:

```
aws amp describe-alert-manager-definition --workspace-id workspace_id --
region region
```

If the status is ACTIVE, your new alert manager definition has taken effect. Until that time, your previous alert manager configuration is still active.

Integrate alerts with Amazon Managed Grafana or open source Grafana

Alert rules that you have created in Alertmanager within Amazon Managed Service for Prometheus can be forwarded and viewed in <u>Amazon Managed Grafana</u> and <u>Grafana</u>, unifying your alert rules and alerts in a single environment. Within Amazon Managed Grafana, you can view your alert rules and the alerts that are generated.

Prerequisites

Before starting to integrate Amazon Managed Service for Prometheus into Amazon Managed Grafana, you must have completed the following prerequisites:

• You must have an existing AWS account and IAM credentials to create Amazon Managed Service for Prometheus and IAM roles programmatically.

For more information about creating an AWS account and IAM credentials, see <u>Set up AWS</u>.

- You must have an Amazon Managed Service for Prometheus workspace, and be ingesting data into it. To set up a new workspace, see <u>Create an Amazon Managed Service for Prometheus</u> <u>workspace</u>. You should also be familiar with the Prometheus concepts such as Alertmanager and Ruler. For more information about these topics, see the <u>Prometheus documentation</u>.
- You have an Alertmanager configuration and a rules file already configured in Amazon Managed Service for Prometheus. For more information about Alertmanager in Amazon Managed Service for Prometheus, see <u>Managing and forwarding alerts in Amazon Managed Service for</u> <u>Prometheus with alert manager</u>. For more information about rules, see <u>Using rules to modify or</u> <u>monitor metrics as they are received</u>.
- You must either have Amazon Managed Grafana set up, or the open source version of Grafana running.
 - If you are using Amazon Managed Grafana, you must be using Grafana alerting. For more information see Migrating legacy dashboard alerts to Grafana alerting.
 - If you are using the open source version of Grafana, you must be running version 9.1 or higher.

🚯 Note

You can use earlier versions of Grafana, but you must <u>enable the unified alerting</u> (Grafana alerting) feature, and you might have to set up a <u>sigv4 proxy</u> to make calls from Grafana to Amazon Managed Service for Prometheus. For more information, see <u>Set up Grafana open source or Grafana Enterprise for use with Amazon Managed</u> Service for Prometheus.

- Amazon Managed Grafana must have the following permissions for your Prometheus resources. You must add them to either the service-managed or customer-managed policies described in https://docs.aws.amazon.com/grafana/latest/userguide/AMG-manage-permissions.html.
 - aps:ListRules
 - aps:ListAlertManagerSilences
 - aps:ListAlertManagerAlerts
 - aps:GetAlertManagerStatus
 - aps:ListAlertManagerAlertGroups
 - aps:PutAlertManagerSilences
 - aps:DeleteAlertManagerSilence

Setting up Amazon Managed Grafana

If you have already set up rules and alerts in your Amazon Managed Service for Prometheus instance, the configuration to use Amazon Managed Grafana as a dashboard for those alerts is done entirely within Amazon Managed Grafana.

To configure Amazon Managed Grafana as your alerts dashboard

- 1. Open the Grafana console for your workspace.
- 2. Under **Configurations**, choose **Data sources**.
- Either create or open your Prometheus data source. If you have not previously set up a Prometheus data source, see <u>Step 2: Add the Prometheus data source in Grafana</u> for more information.
- 4. In the Prometheus data source, select Manage alerts via Alertmanager UI.
- 5. Go back to the **Data sources** interface.
- 6. Create a new Alertmanager data source.
- 7. In the Alertmanager data source configuration page, add the following settings:

• Set Implementation to Prometheus.

- For the URL setting, use the URL for your Prometheus workspace, remove everything after the workspace ID, and append /alertmanager to the end. For example, https://apsworkspaces.us-east1.amazonaws.com/workspaces/ws-example-1234-5678abcd-xyz00000001/alertmanager.
- Under Auth, turn on SigV4Auth. This tells Grafana to use the <u>AWS authentication</u> for the requests.
- Under **SigV4Auth Details**, for **Default Region**, provide the region of your Prometheus instance, for example us-east-1.
- Set the **Default** option to true.
- 8. Choose Save and test.
- 9. Your Amazon Managed Service for Prometheus alerts should now be configured to work with your Grafana instance. Verify that you can see any **Alert rules**, **Alert groups** (including active alerts), and **Silences** from your Amazon Managed Service for Prometheus instance in the Grafana **Alerting** page.

Troubleshoot alert manager with CloudWatch Logs

Using <u>Monitor Amazon Managed Service for Prometheus events with CloudWatch Logs</u>, you can troubleshoot Alert Manager and Ruler related issues. This section contains Alert Manager related troubleshooting topics.

Topics

- Empty content warning
- Non ASCII warning
- Invalid key/value warning
- Message limit warning
- No resource based policy error
- Not authorized to call KMS

Empty content warning

When the log contains the following warning

```
"workspaceId": "ws-abcd1234-ef56-78ab-cd90-1234abcd0000",
"message": {
    "log": "Message has been modified because the content was empty."
    "level": "WARN"
  },
   "component": "alertmanager"
}
```

This means that the Alert manager template resolved the outbound alert to an empty message.

Action to take

Validate your Alert manager template and ensure that you have a valid template for all receiver pathways.

Non ASCII warning

When the log contains the following warning

```
{
    "workspaceId": "ws-abcd1234-ef56-78ab-cd90-1234abcd0000",
    "message": {
        "log": "Subject has been modified because it contains control or non-ASCII
    characters."
        "level": "WARN"
    },
    "component": "alertmanager"
}
```

This means that the subject has non-ASCII characters.

Action to take

Remove references in subject field of your template to the labels that might contain non-ASCII characters.

Invalid key/value warning

When the log contains the following warning

```
"workspaceId": "ws-abcd1234-ef56-78ab-cd90-1234abcd0000",
```

{

```
"message": {
    "log": "MessageAttributes has been removed because of invalid key/value,
numberOfRemovedAttributes=1"
    "level": "WARN"
    },
    "component": "alertmanager"
}
```

This means that some of the message attributes have been removed due to keys/values being invalid.

Action to take

Re-evaluate the templates you are using to populate the message attributes, and ensure it is resolving to a valid SNS message attribute. For more information about validating a message to an Amazon SNS topic, see Validating SNS topic

Message limit warning

When the log contains the following warning

```
{
    "workspaceId": "ws-abcd1234-ef56-78ab-cd90-1234abcd0000",
    "message": {
        "log": "Message has been truncated because it exceeds size limit,
        originSize=266K, truncatedSize=12K"
            "level": "WARN"
        },
        "component": "alertmanager"
}
```

This means that some of the message size is too big.

Action to take

Look at the Alert receiver message template and re-work it to fit within the size limit.

No resource based policy error

When the log contains the following error

```
"workspaceId": "ws-abcd1234-ef56-78ab-cd90-1234abcd0000",
"message": {
    "log": "Notify for alerts failed, AMP is not authorized to perform: SNS:Publish
    on resource: arn:aws:sns:us-west-2:12345:testSnsReceiver because no resource-based
    policy allows the SNS:Publish action"
        "level": "ERROR"
    },
    "component": "alertmanager"
}
```

This means that Amazon Managed Service for Prometheus does not have the permissions to submit the alert to the SNS topic specified.

Action to take

Validate that the access policy on your Amazon SNS topic grants Amazon Managed Service for Prometheus the ability to send SNS messages to the topic. Create an SNS Access Policy giving the service aps.amazonaws.com (Amazon Managed Service for Prometheus) access to your Amazon SNS topic. For more information about SNS Access Policies, see <u>Using the Access Policy</u> <u>Language</u> and <u>Example cases for Amazon SNS access control</u> in the *Amazon Simple Notification Service Developer Guide*.

Not authorized to call KMS

When the log contains the following AWS KMS error

```
{
    "workspaceId": "ws-abcd1234-ef56-78ab-cd90-1234abcd0000",
    "message": {
        "log": "Notify for alerts failed, AMP is not authorized to call KMS",
        "level": "ERROR"
    },
    "component": "alertmanager"
}
```

Action to take

Validate that the key policy of the key used to encrypt the Amazon SNS topic allows the Amazon Managed Service for Prometheus service principal aps.amazonaws.com to perform the following actions: kms:GenerateDataKey*, and kms:Decrypt.For more information, see <u>AWS KMS</u> <u>Permissions for SNS Topic</u>.

Logging and monitoring Amazon Managed Service for Prometheus workspaces

Amazon Managed Service for Prometheus uses Amazon CloudWatch to provide data about its operation. You can use CloudWatch metrics to learn about resource usage and requests to your Amazon Managed Service for Prometheus workspaces. You can turn on CloudWatch Logs support to get logs for events that happen in your workspaces.

The following topics describe using CloudWatch in more detail.

Use CloudWatch metrics to monitor Amazon Managed Service for Prometheus resources

Amazon Managed Service for Prometheus vends usage metrics to CloudWatch. These metrics provide visibility about your workspace utilization. The vended metrics can be found in the AWS/Usage and AWS/Prometheus namespaces in CloudWatch. These metrics are available in CloudWatch for no charge. For more information about usage metrics, see <u>CloudWatch usage metrics</u>.

CloudWatch metric name	Resource name	CloudWatch namespace	Description
ResourceCount [*]	RemoteWri teTPS	AWS/Usage	Remote write operations per second
ResourceCount [*]	QueryMetr icsTPS	AWS/Usage	Query operations per second
ResourceCount	IngestionRate	AWS/Usage	Sample ingestion rate Units: count per second Valid Statistics: Average, Minimum, Maximum, Sum

CloudWatch metric name	Resource name	CloudWatch namespace	Description
ResourceCount	ActiveSeries	AWS/Usage	Number of active series per workspace
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum
ResourceCount	ActiveAlerts	AWS/Usage	Number of active alerts per workspace
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum
ResourceCount	SizeOfAlerts	AWS/Usage	Total size of all alerts in the workspace, in bytes
			Units: bytes
			Valid Statistics: Average, Minimum, Maximum, Sum
ResourceCount	Suppresse dAlerts	AWS/Usage	Number of alerts in suppressed state per workspace. An alert can be suppressed by a silence or inhibition.
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum

CloudWatch metric name	Resource name	CloudWatch namespace	Description
ResourceCount	Unprocess edAlerts	AWS/Usage	Number of alerts in unprocessed state per workspace. An alert is in unprocessed state once it is received by AlertMana ger, but is waiting for the next aggregation group evaluation. Units: count Valid Statistics: Average, Minimum, Maximum, Sum
ResourceCount	AllAlerts	AWS/Usage	Number of alerts in any state per workspace. Units: count Valid Statistics: Average, Minimum, Maximum, Sum
ActiveSer iesPerLabelSet	-	AWS/Prometheus	The current active series usage for each user-defi ned label set Units: count Valid Statistics: Average, Minimum, Maximum, Sum

CloudWatch metric name	Resource name	CloudWatch namespace	Description
ActiveSer iesLimitP erLabelSet	-	AWS/Prometheus	The current active series limit value for each user- defined label set
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum
AlertMana gerAlerts Received	-	AWS/Prometheus	Total successful alerts received by alert manager
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum
AlertMana gerNotifi cationsFailed	-	AWS/Prometheus	Number of failed alert deliveries
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum
AlertMana gerNotifi cationsThrottled	-	AWS/Prometheus	Number of throttled alerts
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum

CloudWatch metric name	Resource name	CloudWatch namespace	Description
Discarded Samples ^{**}	-	AWS/Prometheus	Number of discarded samples by reason
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum
Discarded SamplesPe rLabelSet	-	AWS/Prometheus	The count of discarded samples for each user-defined label set
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum
Ingestion RatePerLabelSet	-	AWS/Prometheus	The ingestion rate for each user-defined label set
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum
QuerySamp lesProcessed	-	AWS/Prometheus	Number of query samples processed
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum

CloudWatch metric name	Resource name	CloudWatch namespace	Description
RuleEvaluations	-	AWS/Prometheus	Total number of rule evaluations
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum
RuleEvalu ationFailures	-	AWS/Prometheus	Number of rule evaluation failures in the interval
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum
RuleGroup IterationsMissed	-	AWS/Prometheus	Number of Rule Group iterations missed in the interval.
			Units: count
			Valid Statistics: Average, Minimum, Maximum, Sum
RuleGroup LastEvalu ationDuration	-	AWS/Prometheus	Duration of a rule group's last evaluation.
			Units: seconds
			Valid Statistics: Average, Minimum, Maximum, Sum

^{*}TPS metrics are generated every minute and are a per-second average over that minute. Short burst periods will not be captured in the TPS metrics.

^{**}Some of the reasons that cause samples to be discarded are as follows.

Reason	Meaning	
greater_than_max_sample_age	Discarding samples which are older than one hour.	
new-value-for-timestamp	Duplicate samples are sent with a different timestamp than was previously recorded.	
per_labelset_series_limit	User has hit the total number of active series per label set limit.	
per_metric_series_limit	User has hit the active series per metric limit.	
per_user_series_limit	User has hit the total number of active series limit.	
rate_limited	Ingestion rate limited.	
sample-out-of-order	Samples are sent out of order and cannot be processed.	
label_value_too_long	Label value is longer than allowed character limit.	
max_label_names_per_series	User has hit the label names per metric.	
missing_metric_name	Metric name is not provided.	
metric_name_invalid	Invalid metric name provided.	
label_invalid	Invalid label provided.	
duplicate_label_names	Duplicate label names provided.	

(i) Note

A metric not existing or missing is the same as the value of that metric being 0.

🚯 Note

RuleGroupIterationsMissed, RuleEvaluations, RuleEvaluationFailures, and RuleGroupLastEvaluationDuration have the RuleGroup dimension of the following structure:

RuleGroupNamespace;RuleGroup

Setting a CloudWatch alarm on Prometheus vended metrics

You can monitor usage of Prometheus resources using CloudWatch alarms.

To set an alarm on the number of ActiveSeries in Prometheus

1. Choose the **Graphed metrics** tab and scroll down to the **ActiveSeries** label.

In the Graphed metrics view, only the metrics currently being ingested will appear.

- 2. Choose the **notification** icon in the **Actions** column.
- 3. In **Specify metric and conditions**, enter the threshold condition in the **Conditions value** field and choose **Next**.
- 4. In **Configure actions**, select an existing SNS topic or create a new SNS topic to send the notification to.
- 5. In **Add name and description**, add the name of the alarm and an optional description.
- 6. Choose Create alarm.

Monitor Amazon Managed Service for Prometheus events with CloudWatch Logs

Amazon Managed Service for Prometheus logs Alert Manager and Ruler error and warning events in log groups in Amazon CloudWatch Logs. For more information about Alert Manager and Rulers, see <u>Alert Manager</u> topic in this guide. You can publish the workspace logs data to log streams in CloudWatch Logs. You can configure the logs that you wish to monitor in the Amazon Managed Service for Prometheus console or by using the AWS CLI. You can view or query these logs in the CloudWatch console. For more information about viewing CloudWatch Logs log streams in the console, see <u>Working with log groups and log streams in CloudWatch</u> in the CloudWatch user guide. The CloudWatch free tier allows up to 5Gb of logs to be published in CloudWatch Logs. The logs that exceed the free tier allowance will be charged based on the CloudWatch pricing plan.

Topics

<u>Configuring CloudWatch Logs</u>

Configuring CloudWatch Logs

Amazon Managed Service for Prometheus logs Alert Manager and Ruler error and warning events in log groups in Amazon CloudWatch Logs.

You can set CloudWatch Logs logging configuration in Amazon Managed Service for Prometheus console or in the AWS CLI by calling the create-logging-configuration API request.

Prerequisites

Before calling create-logging-configuration, attach the following policy or equivalent permissions to the ID or role you will use to configure CloudWatch Logs.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "logs:CreateLogDelivery",
                "logs:GetLogDelivery",
                "logs:UpdateLogDelivery",
                "logs:DeleteLogDelivery",
                "logs:ListLogDeliveries",
                "logs:PutResourcePolicy",
                "logs:DescribeResourcePolicies",
                "logs:DescribeLogGroups",
                "aps:CreateLoggingConfiguration",
                "aps:UpdateLoggingConfiguration",
                "aps:DescribeLoggingConfiguration",
                "aps:DeleteLoggingConfiguration"
            ],
            "Resource": "*"
        }
    ]
```

}

To configure CloudWatch Logs

You can configure logging in Amazon Managed Service for Prometheus using either the AWS console or the AWS CLI.

Console

To configure logging in Amazon Managed Service for Prometheus console

- 1. Navigate to the **Logs** tab in your workspace details panel.
- 2. Choose Manage logs on the upper right side of the Logs panel.
- 3. Choose all in the Log level dropdown list.
- 4. Choose the log group that you want to publish your logs to in the **Log Group** dropdown list.

You can also create a new log group in CloudWatch console.

5. Choose Save changes.

AWS CLI

You can set logging configuration using the AWS CLI.

To configure logging using the AWS CLI

• Using the AWS CLI, run the following command.

Limitations

• Not all events logged

Amazon Managed Service for Prometheus only logs events that are at the warning or error level.

• Policy size limits

CloudWatch Logs resource policies are limited to 5120 characters. When CloudWatch Logs detect that a policy approaches this size limit, it automatically enables log groups that start with /aws/ vendedlogs/.

When you create an alert rule with logging enabled, Amazon Managed Service for Prometheus must update your CloudWatch Logs resource policy with the log group you specify. To avoid reaching the CloudWatch Logs resource policy size limit, prefix your CloudWatch Logs log group names with /aws/vendedlogs/. When you create a log group in the Amazon Managed Service for Prometheus console, the log group names are prefixed with /aws/vendedlogs/. For more information, see Enabling Logging from Certain AWS Services in the CloudWatch Logs User Guide.

Managing the query cost in Amazon Managed Service for Prometheus

Amazon Managed Service for Prometheus offers the ability to limit query cost by providing limits on how much Query Samples Processed (QSP) can be used by a single query. You can configure two types of thresholds for QSP, *warning* and *error* to help manage and control query costs effectively.

When queries hit the *warning* threshold, a warning message appears in the API query response. For queries viewed through Amazon Managed Grafana, the warning will be visible in the Amazon Managed Grafana UI, helping users identify expensive queries. Queries that hit the *error* threshold are not charged and will be rejected with an error.

In addition to query throttling, Amazon Managed Service for Prometheus offers the ability to log query performance data to CloudWatch Logs. This feature allows you to analyze queries in detail, helping you optimize your Amazon Managed Service for Prometheus queries and manage costs more effectively. Query logging captures information about queries that exceed specified Query Samples Processed (QSP) thresholds. This data is then published to CloudWatch Logs, enabling you to investigate and analyze query performance. Logged queries include both API queries and Rule queries. By default, query logging is disabled to minimize unnecessary CloudWatch Logs usage. You can enable this feature when needed for query analysis.

Topics

- Configuring query logging
- Configuring query throttling thresholds

- Log content
- Limitations

Configuring query logging

You can configure query logging in Amazon Managed Service for Prometheus console or in the AWS CLI by calling the create-query-logging-configuration API request. This API body contains list of destinations, but for now, we only support CloudWatch Logs as a destination and destinations should contain exactly one element with CloudWatch configurations.

Prerequisites

Make sure the logGroup is already created. The ID or role used to configure should have the following policy or equivalent permissions.

```
{
"Version": "2012-10-17",
    "Statement": [
        {
"Effect": "Allow",
            "Action": [
                "logs:CreateLogDelivery",
                "logs:GetLogDelivery",
                "logs:UpdateLogDelivery",
                "logs:DeleteLogDelivery",
                "logs:ListLogDeliveries",
                "logs:PutResourcePolicy",
                "logs:DescribeResourcePolicies",
                "logs:DescribeLogGroups",
                "aps:CreateQueryLoggingConfiguration",
                "aps:UpdateQueryLoggingConfiguration",
                "aps:DescribeQueryLoggingConfiguration",
                "aps:DeleteQueryLoggingConfiguration"
            ],
            "Resource": "*"
        }
    ]
}
```

Configure CloudWatch Logs

You can configure CloudWatch Logs by logging into Amazon Managed Service for Prometheus using either the AWS Management Console or the AWS CLI.

To configure query logging using Amazon Managed Service for Prometheus console

- 1. Navigate to the **Logs** tab in your workspace details panel.
- 2. Under **Query Insights**, choose **Create**.
- 3. Select the Log Group drop down and choose the log group to publish your logs.

You can also create a new log group in the CloudWatch console.

- 4. Enter the **Threshold (QSP)**.
- 5. Choose **Save**.

To configure query logging using the AWS CLI use the command

```
aws amp create-query-logging-configuration \
--workspace-id my_workspace_ID \
--destinations '[{"cloudWatchLogs":{"logGroupArn":"$my-log-group-arn"},"filters":
{"qspThreshold":$qspThreshold}}]'
```

For information on how to update, delete, and describe operations, see <u>Amazon Managed Service</u> for Prometheus API Reference.

Configuring query throttling thresholds

To configure QSP thresholds, you must provide the query parameters in the QueryMetrics API.

- max_samples_processed_warning_threshold Sets the warning threshold for query samples processed
- max_samples_processed_error_threshold Sets the error threshold for query samples processed

For Amazon Managed Grafana users, you can use grafana data source configuration to apply limits to all the queries from the datasource:

1. Browse to the Amazon Managed Service for Prometheus data source configuration in Amazon Managed Grafana.

- 2. Under **Custom query parameters**, add the threshold headers.
- 3. Choose Save.

Log content

For queries that originate from rules, you will see the following information about the query in the CloudWatch Logs:

```
{
    workspaceId: "workspace_id",
    message: {
        query: "avg(rate(go_goroutines[1m])) > 1",
        name: "alert_rule",
        kind: "alerting",
        group: "test-alert",
        namespace: "test",
        samples: "59321",
    },
    component: "ruler"
}
```

For queries that originate from API calls, you will see the following information about the query in the CloudWatch Logs:

```
{
    workspaceId: "ws-5e7658c2-7ccf-4c30-9de9-2ab26fa30639",
    message: {
        query: "sum by (instance) (go_memstats_alloc_bytes{job=\"node\"})",
        queryType: "range",
        start: "1683308700000",
        end: "1683913500000",
        step: "300000",
        step: "300000",
        samples: "11496",
        userAgent: "AWSPrometheusDPJavaClient/2.0.436.0 ",
        dashboardUid: "11234",
        paneIId: "12"
    },
    component: "query-frontend"
}
```

Limitations

Policy size limits – CloudWatch Logs resource policies are limited to 5120 characters. When CloudWatch Logs detects that the policy is approaching the size limit, it automatically enables log groups that start with /aws/vendedlogs/. When you enable query logging, Amazon Managed Service for Prometheus must update your CloudWatch Logs resource policy with the log group you specify. To avoid reaching the CloudWatch Logs resource policy size limit, prefix your CloudWatch Logs log group names with /aws/vendedlogs/.

Understand and optimize costs in Amazon Managed Service for Prometheus

The following frequently asked questions and their answers may be helpful in understanding and optimizing costs associated with Amazon Managed Service for Prometheus.

What contributes to my costs?

For most customers, metric *ingestion* contributes the majority of costs. Customers with high query usage will also see some cost based on *query samples processed*, with *metrics storage* being a small driver of overall costs. For more information about the prices for each of these, see <u>Pricing</u> in the *Amazon Managed Service for Prometheus product page*.

What is the best way to lower my costs? How do I lower ingestion costs?

Ingestion rates (not storage of the metrics) is the majority of costs for most customers. You can reduce ingestion rates by reducing the collection frequency (increasing the collection interval) or by reducing the number of active series ingested.

You can increase the collection (scraping) interval from your collection agent: Both the Prometheus server (running in Agent mode) and the AWS Distro for OpenTelemetry (ADOT) collector support the scrape_interval configuration. For example, increasing the collection interval from 30 seconds to 60 seconds will reduce your ingestion usage by half.

You can also filter the metrics sent to Amazon Managed Service for Prometheus by using the <relabel_config>. For more information about relabeling in the Prometheus agent configuration, see https://prometheus.io/docs/prometheus/latest/configuration/configuration/ #relabel_config in the Prometheus documentation.

What is the best way to lower my query costs?

Query costs are based on the number of samples processed. You can reduce the frequency of queries to reduce your query costs.

To get more visibility into the queries that are contributing the most to your query costs, you can reach out to file a ticket with your support contact. The Amazon Managed Service for Prometheus team can help you understand the queries that are contributing the most to your costs.

If I decrease the retention period of my metrics, will that help reduce my total bill?

You can reduce your retention period, however, this is unlikely to substantially reduce your costs.

For information about how to configure the retention period of a workspace, see <u>Configure your</u> <u>workspace</u>.

How can I keep my alert query costs low?

Alerting creates queries against your data, which add to your query costs. Here are some strategies that you can use to optimize your alert queries, and keep your costs lower.

 Use Amazon Managed Service for Prometheus alerting – Alerting systems external to Amazon Managed Service for Prometheus may require additional queries to add resiliency or high availability, as the external service queries the metrics from multiple availability zones or regions. This includes alerting in Grafana for high availability. This can multiply your cost by three times or more. The alerting in Amazon Managed Service for Prometheus is optimized and will give you high availability and resiliency with the fewest number of queries.

We recommend using the native alerting in Amazon Managed Service for Prometheus rather than external alerting systems.

- Optimize your alert interval One quick way to optimize your alert queries is to increase the auto-refresh interval. If you have an alert that queries every minute, but is only needed every five minutes, increasing the auto-refresh interval could save you five times your query costs for that alert.
- Use an optimal lookback A larger lookback window in your query increases the costs of the query, as it pulls more data. Ensure that the lookback window in your PromQL query is reasonably sized for the data you need to alert. For example, in the following rule, the expression includes a ten minute lookback window:

```
- alert: metric:alerting_rule
  expr: avg(rate(container_cpu_usage_seconds_total[10m])) > 0
```

for: 2m

Changing the expr to avg(rate(container_cpu_usage_seconds_total[5m])) > 0 can help to reduce your query costs.

In general, look at your alerting rules and make sure that you are alerting on the best metrics for your service. It's easy to create overlapping alerts on the same metrics or multiple alerts that give you the same information, especially as you add alerts over time. If you find that you often see groups of alerts happening at the same time, it's possible that you can optimize your alerts and not include all of them.

These suggestions can help you to reduce costs. Ultimately, you must balance the costs with creating the right set of alerts for understanding the state of your system.

For more information about alerting in Amazon Managed Service for Prometheus, see <u>Managing</u> and forwarding alerts in Amazon Managed Service for Prometheus with alert manager.

What metrics can I use to monitor my costs?

Monitor IngestionRate in Amazon CloudWatch to track your ingestion costs. For more information about monitoring Amazon Managed Service for Prometheus metrics in CloudWatch, see <u>Use CloudWatch metrics to monitor Amazon Managed Service for Prometheus resources</u>.

Can I check my bill at any time?

The AWS Cost and Usage Report tracks your AWS usage and provides estimated charges associated with your account within a billing period. For more information, see <u>What are AWS Cost and Usage</u> <u>Reports?</u> in the AWS Cost and Usage Reports User Guide

Why is my bill higher at the beginning of the month than at the end of the month?

Amazon Managed Service for Prometheus has a tiered pricing model for ingestion, which results in costs in your initial usage being higher. As your usage reaches higher ingest tiers, with lower costs, your costs are lower. For more information about pricing, including ingest tiers, see <u>Pricing</u> in the *Amazon Managed Service for Prometheus product page*.

🚯 Note

- Tiers are for usage *within a region*, not across regions. Usage within a region must reach the next tier to use the lower rate.
- In an organization in AWS Organizations, tier usage is tallied *per payer account*, not per account (the payer account is always the organization management account). When the total ingested metrics (within a region) for *all accounts in an organization* reaches the next tier, all accounts are charged the lower rate.

I deleted all my Amazon Managed Service for Prometheus workspaces, but I still seem to be getting charged. What might be happening?

One possibility in this case is that you still have AWS managed scrapers that are setup to send metrics to your deleted workspaces. Follow the instructions to <u>Find and delete scrapers</u>.

Integrating with other AWS services

Amazon Managed Service for Prometheus integrates with other AWS services. This section describes integrating with Amazon Elastic Kubernetes Service (Amazon EKS) cost monitoring (with Kubecost), and how to ingest metrics from CloudWatch using Amazon Data Firehose. It also describes setting up and managing Amazon Managed Service for Prometheus with AWS Observability Accelerator Terraform modules, or by using AWS Controllers for Kubernetes.

Topics

- Integrating with Amazon EKS cost monitoring
- Set up Amazon Managed Service for Prometheus with AWS Observability Accelerator
- Manage Amazon Managed Service for Prometheus with AWS Controllers for Kubernetes
- Integrating CloudWatch metrics with Amazon Managed Service for Prometheus

Integrating with Amazon EKS cost monitoring

Amazon Managed Service for Prometheus integrates with Amazon Elastic Kubernetes Service (Amazon EKS) cost monitoring (with Kubecost) to perform cost allocation calculations and provide insights into optimizing your Kubernetes clusters. Using Amazon Managed Service for Prometheus with Kubecost, you can reliably scale your cost monitoring to support larger clusters.

Integrating with Kubecost gives you granular visibility into your Amazon EKS cluster costs. You can aggregate costs by the majority of Kubernetes contexts, from the container level up to the cluster level, and even multi-cluster level. You can generate reports across containers or clusters to track costs for show back or chargeback purposes.

The following give instructions for integrating with Kubecost in a single- or multi-cluster scenario:

- Single-cluster integration To learn how to integrate Amazon EKS cost monitoring with a single cluster, see the AWS blog post <u>Integrating Kubecost with Amazon Managed Service for</u> <u>Prometheus</u>.
- Multi-cluster integration To learn how to integrate Amazon EKS cost monitoring with a multiple clusters, see the AWS blog post <u>Multi-cluster cost monitoring for Amazon EKS using</u> Kubecost and Amazon Managed Service for Prometheus.

For more information about using Kubecost, see <u>Cost monitoring</u> in the *Amazon EKS User Guide*.

Set up Amazon Managed Service for Prometheus with AWS Observability Accelerator

AWS provides observability tools, including monitoring, logging, alerting, and dashboards, for your Amazon Elastic Kubernetes Service (Amazon EKS) projects. This includes Amazon Managed Service for Prometheus, <u>Amazon Managed Grafana</u>, <u>AWS Distro for OpenTelemetry</u>, and other tools. To help you use these tools together, AWS provides Terraform modules that configure observability with these services, called the AWS Observability Accelerator.

AWS Observability Accelerator provides examples for monitoring infrastructure, <u>NGINX</u> deployements, and other scenarios. This section gives an example of monitoring infrastructure within your Amazon EKS cluster.

The Terraform templates and detailed instructions can be found on the <u>AWS Observability</u> <u>Accelerator for Terraform GitHub page</u>. You can also read the <u>blog post announcing AWS</u> Observability Accelerator.

Prerequisites

To use AWS Observability Accelerator, you must have an existing Amazon EKS cluster, and the following prerequisites:

- AWS CLI used to call AWS functionality from the command line.
- <u>kubectl</u> used to control your EKS cluster from the command line.
- <u>Terraform</u> used to automate creation of the resources for this solution. You must have the AWS provider setup with an IAM role that has access to create and manage Amazon Managed Service for Prometheus, Amazon Managed Grafana, and IAM within your AWS account. For more information about how to configure the AWS provider for Terraform, see <u>AWS provider</u> in the *Terraform documentation*.

Using the infrastructure monitoring example

AWS Observability Accelerator provides example templates that use the included Terraform modules to set up and configure observability for your Amazon EKS cluster. This example demonstrates using AWS Observability Accelerator to set up infrastructure monitoring. For more details about using this template and additional capabilities that it includes, see <u>Existing Cluster</u> with the AWS Observability Accelerator base and Infrastructure monitoring page on GitHub.

To use the infrastructure monitoring Terraform module

1. From the folder you want to create your project in, clone the repo using the following command.

git clone https://github.com/aws-observability/terraform-aws-observabilityaccelerator.git

2. Initialize Terraform with the following commands.

cd examples/existing-cluster-with-base-and-infra

terraform init

3. Create a new terraform.tfvars file, as in the following example. Use the AWS Region and cluster ID for your Amazon EKS cluster.

```
# (mandatory) AWS Region where your resources will be located
aws_region = "eu-west-1"
# (mandatory) EKS Cluster name
eks_cluster_id = "my-eks-cluster"
```

- 4. Create an Amazon Managed Grafana workspace, if you don't already have one that you want to use. For information about how to create a new workspace, see <u>Create your first workspace</u> in the *Amazon Managed Grafana User Guide*.
- 5. Create two variables for Terraform to use your Grafana workspace by running the following commands at the command line. You will need to replace the *grafana-workspace-id* with the ID from your Grafana workspace.

export TF_VAR_managed_grafana_workspace_id=grafana-workspace-id

```
export TF_VAR_grafana_api_key=`aws grafana create-workspace-api-key --key-name
"observability-accelerator-$(date +%s)" --key-role ADMIN --seconds-to-live 1200 --
workspace-id $TF_VAR_managed_grafana_workspace_id --query key --output text`
```

6. [Optional] To use an existing Amazon Managed Service for Prometheus workspace, add the ID to the terraform.tfvars file, as in the following example, replacing the *prometheus-workspace-id* with your Prometheus workspace ID. If you do not specify an existing workspace, then a new Prometheus workspace will be created for you.

```
# (optional) Leave it empty for a new workspace to be created
managed_prometheus_workspace_id = "prometheus-workspace-id"
```

7. Deploy the solution with the following command.

```
terraform apply -var-file=terraform.tfvars
```

This will create resources in your AWS account, including the following:

- A new Amazon Managed Service for Prometheus workspace (unless you opted to use an existing workspace).
- Alert manager configuration, alerts, and rules in your Prometheus workspace.
- New Amazon Managed Grafana data source and dashboards in your current workspace. The data source will be called aws-observability-accelerator. The dashboards will be listed under Observability Accelerator Dashboards.
- An <u>AWS Distro for OpenTelemetry</u> operator set up in the provided Amazon EKS cluster, to send metrics to your Amazon Managed Service for Prometheus workspace.

To view your new dashboards, open the specific dashboard in your Amazon Managed Grafana workspace. For more information about using Amazon Managed Grafana, see <u>Working in your</u> Grafana workspace, in the Amazon Managed Grafana User Guide.

Manage Amazon Managed Service for Prometheus with AWS Controllers for Kubernetes

Amazon Managed Service for Prometheus is integrated with <u>AWS Controllers for Kubernetes</u> (ACK), with support for managing your workspace, Alert Manager, and Ruler resources in Amazon EKS. You can use AWS Controllers for Kubernetes custom resource definitions (CRDs) and native Kubernetes objects without having to define any resources outside of your cluster.

This section describes how to set up AWS Controllers for Kubernetes and Amazon Managed Service for Prometheus in an existing Amazon EKS cluster.

You can also read the blog posts <u>introducing AWS Controllers for Kubernetes</u> and <u>introducing the</u> ACK controller for Amazon Managed Service for Prometheus.

Prerequisites

Before starting to integrate AWS Controllers for Kubernetes and Amazon Managed Service for Prometheus with your Amazon EKS cluster, you must have the following prerequisites.

- You must have an <u>existing AWS account and permissions</u> to create Amazon Managed Service for Prometheus and IAM roles programmatically.
- You must have an existing <u>Amazon EKS cluster</u> with OpenID Connect (OIDC) enabled.

If you do not have OIDC enabled, you can use the following command to enable it. Remember to replace the *YOUR_CLUSTER_NAME* and *AWS_REGION* with the correct values for your account.

```
eksctl utils associate-iam-oidc-provider \
    --cluster ${YOUR_CLUSTER_NAME} --region ${AWS_REGION} \
    --approve
```

For more information about using OIDC with Amazon EKS, see <u>OIDC identity provider</u> authentication and <u>Creating an IAM OIDC provider</u> in the *Amazon EKS User Guide*.

- You must have the Amazon EBS CSI driver installed in your Amazon EKS cluster.
- You must have the <u>AWS CLI</u> installed. The AWS CLI is used to call AWS functionality from the command line.
- Helm, the package manager for Kubernetes, must be installed.
- Control plane metrics with Prometheus must be set up in your Amazon EKS cluster.
- You must have an <u>Amazon Simple Notification Service (Amazon SNS)</u> topic where you want to send alerts from your new workspace. Make sure that you have <u>given Amazon Managed Service</u> for Prometheus permission to send messages to the topic.

When your Amazon EKS cluster is configured appropriately, you should be able to see metrics formatted for Prometheus by calling kubectl get --raw /metrics. Now you are ready to install an AWS Controllers for Kubernetes service controller and use it to deploy Amazon Managed Service for Prometheus resources.

Deploying a workspace with AWS Controllers for Kubernetes

To deploy a new Amazon Managed Service for Prometheus workspace, you will install an AWS Controllers for Kubernetes controller, and then use that to create the workspace.

To deploy a new Amazon Managed Service for Prometheus workspace with AWS Controllers for Kubernetes

 Use the following commands to use Helm to install the Amazon Managed Service for Prometheus service controller. For more information see <u>Install an ACK Controller</u> in the AWS Controllers for Kubernetes documentation on GitHub. Use the correct *region* for your system, such as us-east-1.

```
export SERVICE=prometheusservice
export RELEASE_VERSION=`curl -sL https://api.github.com/repos/aws-controllers-k8s/
$SERVICE-controller/releases/latest | grep '"tag_name":' | cut -d'"' -f4`
export ACK_SYSTEM_NAMESPACE=ack-system
export AWS_REGION=region
aws ecr-public get-login-password --region us-east-1 | helm registry login --
username AWS --password-stdin public.ecr.aws
helm install --create-namespace -n $ACK_SYSTEM_NAMESPACE ack-$SERVICE-controller \
oci://public.ecr.aws/aws-controllers-k8s/$SERVICE-chart --version=
$RELEASE_VERSION --set=aws.region=$AWS_REGION
```

After a few moments, you should see a response similar to the following indicating success.

You are now able to create Amazon Managed Service for Prometheus (AMP) resources! The controller is running in "cluster" mode. The controller is configured to manage AWS resources in region: "us-east-1"

You can optionally verify that the AWS Controllers for Kubernetes controller has been successfully installed with the following command.

helm list --namespace \$ACK_SYSTEM_NAMESPACE -o yaml

This will return information about the controller ack-prometheusservice-controller, including the status: deployed.

2. Create a file called workspace.yaml with the following text. This will be used as configuration for the workspace you are creating.

```
apiVersion: prometheusservice.services.k8s.aws/v1alpha1
kind: Workspace
metadata:
   name: my-amp-workspace
spec:
   alias: my-amp-workspace
   tags:
      ClusterName: EKS-demo
```

3. Run the following command to create your workspace (this command depends on the system variables that you set up in step 1).

kubectl apply -f workspace.yaml -n \$ACK_SYSTEM_NAMESPACE

Within a few moments, you should be able to see a new workspace, called my-amp-workspace in your account.

Running the following command to view the details and status of your workspace including the *workspace ID*. Alternately, you can view the new workspace in the <u>Amazon Managed</u> <u>Service for Prometheus console</u>.

```
kubectl describe workspace my-amp-workspace -n $ACK_SYSTEM_NAMESPACE
```

🚯 Note

You can also use an existing workspace rather than create a new one.

4. Create two new yaml files as configuration for the Rulegroups and AlertManager that you will create next using the following configuration.

Save this configuration as rulegroup.yaml. Replace *WORKSPACE-ID* with the workspace ID from the previous step.

```
apiVersion: prometheusservice.services.k8s.aws/v1alpha1
kind: RuleGroupsNamespace
metadata:
  name: default-rule
spec:
  workspaceID: WORKSPACE-ID
  name: default-rule
  configuration: |
    groups:
    - name: example
      rules:
      - alert: HostHighCpuLoad
        expr: 100 - (avg(rate(node_cpu_seconds_total{mode="idle"}[2m])) * 100) > 60
        for: 5m
        labels:
          severity: warning
          event_type: scale_up
        annotations:
          summary: Host high CPU load (instance {{ $labels.instance }})
          description: "CPU load is > 60%\n VALUE = {{ $value }}\n LABELS =
 {{ $labels }}"
      - alert: HostLowCpuLoad
        expr: 100 - (avg(rate(node_cpu_seconds_total{mode="idle"}[2m])) * 100) < 30</pre>
        for: 5m
        labels:
          severity: warning
          event_type: scale_down
        annotations:
          summary: Host low CPU load (instance {{ $labels.instance }})
          description: "CPU load is < 30%\n VALUE = {{ $value }}\n LABELS =</pre>
 {{ $labels }}"
```

Save the following configuration as alertmanager.yaml. Replace *WORKSPACE-ID* with the workspace ID from the previous step. Replace *TOPIC-ARN* with the ARN for the Amazon SNS topic to send notifications to, and *REGION* with the AWS Region you are using. Remember that Amazon Managed Service for Prometheus <u>must have permissions</u> to the Amazon SNS topic.

```
apiVersion: prometheusservice.services.k8s.aws/v1alpha1
```

```
kind: AlertManagerDefinition
metadata:
  name: alert-manager
spec:
  workspaceID: WORKSPACE-ID
  configuration: |
    alertmanager_config: |
      route:
         receiver: default_receiver
      receivers:
        - name: default_receiver
          sns_configs:
          - topic_arn: TOPIC-ARN
            siqv4:
              region: REGION
            message: |
              alert_type: {{ .CommonLabels.alertname }}
              event_type: {{ .CommonLabels.event_type }}
```

🚯 Note

To learn more about the formats of these configuration files, see <u>RuleGroupsNamespaceData</u> and <u>AlertManagerDefinitionData</u>.

5. Run the following commands to create your rule group and alert manager configuration (this command depends on the system variables that you set up in step 1).

kubectl apply -f rulegroup.yaml -n \$ACK_SYSTEM_NAMESPACE kubectl apply -f alertmanager.yaml -n \$ACK_SYSTEM_NAMESPACE

The changes will be available within a few moments.

🚯 Note

To update a resource, rather than create it, you simply update the yaml file, and run the kubectl apply command again.

To delete a resource, run the following command. Replace *ResourceType* with the type of resource you want to delete Workspace, AlertManagerDefinition, or RuleGroupNamespace. Replace *ResourceName* with the name of the resource to delete.

User Guide

kubectl delete ResourceType ResourceName -n \$ACK_SYSTEM_NAMESPACE

That completes deploying the new workspace. The next section describes configuring your cluster to send metrics to that workspace.

Configuring your Amazon EKS cluster to write to the Amazon Managed Service for Prometheus workspace

This section describes how to use Helm to configure the Prometheus running in your Amazon EKS cluster to remote write metrics to the Amazon Managed Service for Prometheus workspace that you created in the previous section.

For this procedure, you will need the name of the IAM role you have created to use for ingesting metrics. If you have not done this already, see <u>Set up service roles for the ingestion of metrics from</u> <u>Amazon EKS clusters</u> for more information and instructions. If you follow those instructions, the IAM role will be called amp-iamproxy-ingest-role.

To configure your Amazon EKS cluster for remote write

 Use the following command to get the prometheusEndpoint for your workspace. Replace WORKSPACE-ID with the workspace ID from the previous section.

aws amp describe-workspace --workspace-id WORKSPACE-ID

The prometheusEndpoint will be in the return results, and be formatted like this:

```
https://aps-workspaces.us-west-2.amazonaws.com/workspaces/ws-a1b2c3d4-a123-b456-
c789-ac1234567890/
```

Save this URL for use in the next few steps.

 Create a new file with the following text and call it prometheus-config.yaml. Replace account with your account ID, workspaceURL/ with the URL you just found, and region with the appropriate AWS Region for your system.

```
serviceAccounts:
server:
```

3. Find the Prometheus chart and namespace names as well as the chart version with the following Helm command.

helm ls --all-namespaces

Based on the steps so far, the Prometheus chart and namespace should both be named prometheus, and the chart version may be 15.2.0

4. Run the following command, using the *PrometheusChartName*, *PrometheusNamespace*, and *PrometheusChartVersion* found in the previous step.

helm upgrade PrometheusChartName prometheus-community/prometheus n PrometheusNamespace -f prometheus-config.yaml --version PrometheusChartVersion

After a few minutes, you'll see a message that the upgrade was successful.

5. Optionally, validate that metrics are succesfully being sent by querying the Amazon Managed Service for Prometheus endpoint via awscurl. Replace *Region* with the AWS Region that you are using, and *workspaceURL/* with the URL you found in step 1.

```
awscurl --service="aps" --region="Region" "workspaceURL/api/v1/query?
query=node_cpu_seconds_total"
```

You have now created an Amazon Managed Service for Prometheus workspace and connected to it from your Amazon EKS cluster, using YAML files as configuration. These files, called custom resource definitions (CRDs), live within your Amazon EKS cluster. You can use the AWS Controllers for Kubernetes controller to manage all of your Amazon Managed Service for Prometheus resources directly from the cluster.

Integrating CloudWatch metrics with Amazon Managed Service for Prometheus

It can help to have all your metrics in one place. Amazon Managed Service for Prometheus does not automatically ingest Amazon CloudWatch metrics. However, you can use Amazon Data Firehose and AWS Lambda to push CloudWatch metrics to Amazon Managed Service for Prometheus.

This section describes how to instrument a <u>Amazon CloudWatch metric stream</u> and use <u>Amazon</u> <u>Data Firehose</u> and <u>AWS Lambda</u> to ingest metrics into Amazon Managed Service for Prometheus.

You will set up a stack using <u>AWS Cloud Development Kit (CDK)</u> to create a Firehose Delivery Stream, a Lambda, and an Amazon S3 bucket to demonstrate a complete scenario.

Infrastructure

The first thing you must do is set up the infrastructure for this recipe.

CloudWatch metric streams allow forwarding of the streaming metric data to an HTTP endpoint or <u>Amazon S3 bucket</u>.

Setting up the infrastructure will consist of 4 steps:

- Configuring prerequisites
- Creating an Amazon Managed Service for Prometheus workspace
- Installing dependencies
- Deploying the stack

Prerequisites

- The AWS CLI is installed and configured in your environment.
- The AWS CDK Typescript is installed in your environment.
- Node.js and Go are installed in your environment.
- The <u>AWS observability CloudWatch metrics exporter github repository</u> (CWMetricsStreamExporter) has been cloned to your local machine.

 The demo application in this recipe will be running on top of Amazon Managed Service for Prometheus. Create your Amazon Managed Service for Prometheus Workspace via the following command:

aws amp create-workspace --alias prometheus-demo-recipe

2. Ensure your workspace has been created with the following command:

aws amp list-workspaces

For more information about Amazon Managed Service for Prometheus, see <u>Amazon Managed</u> Service for Prometheus User Guide.

To install dependencies

1. Install dependencies

From the root of the aws-olly-recipes repository, change your directory to CWMetricStreamExporter using the command:

cd sandbox/CWMetricStreamExporter

This will now be considered the root of the repo, going forward.

2. Change directory to /cdk via the following command:

cd cdk

3. Install the CDK dependencies via the following command:

npm install

4. Change directory back to the root of the repo, and then change directory to /lambda using the following command:

cd lambda

5. Once in the /lambda folder, install the Go dependencies using:

go get

All the dependencies are now installed.

To deploy the stack

 In the root of the repo, open config.yaml and modify the Amazon Managed Service for Prometheus workspace URL by replacing the {workspace} with the newly created workspace id, and the region your Amazon Managed Service for Prometheus workspace is in.

For example, modify the following with:

```
AMP:
    remote_write_url: "https://aps-workspaces.us-east-2.amazonaws.com/workspaces/
{workspaceId}/api/v1/remote_write"
    region: us-east-2
```

Change the names of the Firehose delivery stream and Amazon S3 bucket to your liking.

2. To build the AWS CDK and the Lambda code, in the root of the repo run the following commend:

npm run build

This build step ensures that the Go Lambda binary is built, and deploys the CDK to CloudFormation.

- 3. To complete the deployment, review and accept the IAM changes that the stack requires.
- 4. (Optional) You can very if that the stack has been created by running the following command.

```
aws cloudformation list-stacks
```

A stack named CDK Stack will be in the list.

Creating a Amazon CloudWatch stream

Now that you have a lambda function to handle the metrics, you can create the metrics stream from Amazon CloudWatch.

To create an CloudWatch metrics stream

- 1. Navigate to the CloudWatch console, at https://console.aws.amazon.com/cloudwatch/ home#metric-streamsList and select Create metric stream.
- 2. Select the metrics needed, either all metrics, or only from selected namespaces.
- 3. Under Configuration, choose Select an existing Firehose owned by your account.
- 4. You will be using the Firehose created earlier by the CDK. In the Select your Kinesis data Firehose stream drop down, select the stream created earlier. It will have a name like CdkStack-KinesisFirehoseStream123456AB-sample1234.
- 5. Change the output format to **JSON**.
- 6. Give the metric stream a name that is meaningful to you.
- 7. Choose Create metric stream.
- 8. (Optional) To verify the Lambda function invocation, navigate to the <u>Lambda console</u> and choose the function KinesisMessageHandler. Select the **Monitor** tab and **Logs** subtab, and under **Recent Invocations** there should be entries of the Lambda function being triggered.

🚯 Note

It may take up to 5 minutes before invocations begin to show in the Monitor tab.

Your metrics are now being streamed from Amazon CloudWatch to Amazon Managed Service for Prometheus.

Cleanup

You may want to clean up the resources that were used in this example. The following procedure explains how to do so. This will stop the metrics stream that you created.

To clean up resources

1. Start by deleting the CloudFormation stack with the following commands:

```
cd cdk
cdk destroy
```

2. Remove the Amazon Managed Service for Prometheus workspace:

```
aws amp delete-workspace --workspace-id \
    `aws amp list-workspaces --alias prometheus-sample-app --query
'workspaces[0].workspaceId' --output text`
```

3. Finally, remove the Amazon CloudWatch metric stream using the Amazon CloudWatch console.

Security in Amazon Managed Service for Prometheus

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 Amazon Managed Service for Prometheus, see <u>AWS Services in Scope by Compliance Program</u>.
- 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 Amazon Managed Service for Prometheus. The following topics show you how to configure Amazon Managed Service for Prometheus to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Amazon Managed Service for Prometheus resources.

Topics

- Data protection in Amazon Managed Service for Prometheus
- Identity and Access Management for Amazon Managed Service for Prometheus
- IAM permissions and policies
- <u>Compliance Validation for Amazon Managed Service for Prometheus</u>
- <u>Resilience in Amazon Managed Service for Prometheus</u>
- Infrastructure Security in Amazon Managed Service for Prometheus
- Using service-linked roles for Amazon Managed Service for Prometheus
- Logging Amazon Managed Service for Prometheus API calls using AWS CloudTrail
- <u>Set up IAM roles for service accounts</u>

• Using Amazon Managed Service for Prometheus with interface VPC endpoints

Data protection in Amazon Managed Service for Prometheus

The AWS <u>shared responsibility model</u> applies to data protection in Amazon Managed Service for Prometheus. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the 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 <u>Data Privacy FAQ</u>. For information about data protection in Europe, see the <u>AWS Shared</u> <u>Responsibility Model and GDPR</u> blog post on 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. For information about using CloudTrail trails to capture AWS activities, see <u>Working with CloudTrail trails</u> in the AWS CloudTrail User Guide.
- 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 140-3 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 <u>Federal Information Processing Standard (FIPS) 140-3</u>.

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 Amazon Managed Service for Prometheus 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.

Topics

- Data collected by Amazon Managed Service for Prometheus
- Encryption at rest

Data collected by Amazon Managed Service for Prometheus

Amazon Managed Service for Prometheus collects and stores operational metrics that you configure to be sent from Prometheus servers running in your account to Amazon Managed Service for Prometheus. This data includes the following:

- Metric values
- Metric labels (or arbitrary key-value pairs) that help identify and classify data
- Timestamps for data samples

Unique tenant IDs isolate data from different customers. These IDs limit what customer data is accessible. Customers can't change tenant IDs.

Amazon Managed Service for Prometheus encrypts the data that it stores with AWS Key Management Service (AWS KMS) keys. Amazon Managed Service for Prometheus manages these keys.

Note

Amazon Managed Service for Prometheus supports the creation of customer managed keys for encrypting your data. For more information about the keys that Amazon Managed Service for Prometheus uses by default, and how to use your own customer managed keys, see Encryption at rest.

Data in transit is encrypted with HTTPS automatically. Amazon Managed Service for Prometheus secures connections between Availability Zones within an AWS Region using HTTPS internally.

Encryption at rest

By default, Amazon Managed Service for Prometheus automatically provides you with encryption at rest and does this using AWS owned encryption keys.

 AWS owned keys – Amazon Managed Service for Prometheus uses these keys to automatically encrypt data uploaded to your workspace. You can't view, manage or use AWS owned keys, or audit their use. However, you don't have to take any action or change any programs to protect the keys that encrypt your data. For more information, see <u>AWS owned keys</u> in the AWS Key Management Service Developer Guide.

Encryption of data at rest helps reduce the operational overhead and complexity that goes into protecting sensitive customer data, such as personally identifiable information. It allows you to build secure applications that meet strict encryption compliance and regulatory requirements.

You can alternatively choose to use a customer managed key when you create your workspace:

- Customer managed keys Amazon Managed Service for Prometheus supports the use of a symmetric customer managed key that you create, own, and manage to encrypt the data in your workspace. Because you have full control of this encryption, you can perform such tasks as:
 - Establishing and maintaining key policies
 - Establishing and maintaining IAM policies and grants
 - Enabling and disabling key policies
 - Rotating key cryptographic material
 - Adding tags
 - Creating key aliases
 - Scheduling keys for deletion

For more information, see <u>customer managed keys</u> in the AWS Key Management Service Developer Guide.

Choose whether to use customer managed keys or AWS owned keys carefully. Workspaces created with customer managed keys can't be converted to use AWS owned keys later (and vice versa).

Note

Amazon Managed Service for Prometheus automatically enables encryption at rest using AWS owned keys to protect your data at no charge.

However, AWS KMS charges apply for using a customer managed key. For more information about pricing, see AWS Key Management Service pricing.

For more information on AWS KMS, see What is AWS Key Management Service?

í) Note

Workspaces created with customer managed keys cannot use <u>AWS managed collectors</u> for ingestion.

How Amazon Managed Service for Prometheus uses grants in AWS KMS

Amazon Managed Service for Prometheus requires three grants to use your customer managed key.

When you create an Amazon Managed Service for Prometheus workspace encrypted with a customer managed key, Amazon Managed Service for Prometheus creates the three grants on your behalf by sending <u>CreateGrant</u> requests to AWS KMS. Grants in AWS KMS are used to give Amazon Managed Service for Prometheus access to the KMS key in your account, even when not called directly on your behalf (for example, when storing metrics data that has been scraped from an Amazon EKS cluster.

Amazon Managed Service for Prometheus requires the grants to use your customer managed key for the following internal operations:

- Send <u>DescribeKey</u> requests to AWS KMS to verify that the symmetric customer managed KMS key given when creating a workspace is valid.
- Send <u>GenerateDataKey</u> requests to AWS KMS to generate data keys encrypted by your customer managed key.
- Send <u>Decrypt</u> requests to AWS KMS to decrypt the encrypted data keys so that they can be used to encrypt your data.

Amazon Managed Service for Prometheus creates three grants to the AWS KMS key that allow Amazon Managed Service for Prometheus to use the key on your behalf. You can remove access to the key by changing the key policy, by disabling the key, or by revoking the grant. You should understand the consequences of these actions before performing them. This can cause data loss in your workspace.

If you remove access to any of the grants in any way, Amazon Managed Service for Prometheus won't be able to access any of the data encrypted by the customer managed key, nor store new

data sent to the workspace, which affects operations that are dependent on that data. New data sent to the workspace will not be accessible and may be permanently lost.

🔥 Warning

• If you disable the key, or remove Amazon Managed Service for Prometheus access in the key policy, the workspace data is no longer accessible. New data being sent to the workspace will not be accessible and may be permanently lost.

You can get access to the workspace data and start receiving new data again by restoring Amazon Managed Service for Prometheus access to the key.

• If you *revoke* a grant, it can't be recreated, and the data in the workspace is lost permanently.

Step 1: Create a customer managed key

You can create a symmetric customer managed key by using the AWS Management Console, or the AWS KMS APIs. The key does not need to be in the same account as the Amazon Managed Service for Prometheus workspace, as long as you provide the correct access through policy, as described below.

To create a symmetric customer managed key

Follow the steps for <u>Creating symmetric customer managed key</u> in the AWS Key Management Service Developer Guide.

Key policy

Key policies control access to your customer managed key. Every customer managed key must have exactly one key policy, which contains statements that determine who can use the key and how they can use it. When you create your customer managed key, you can specify a key policy. For more information, see <u>Managing access to customer managed keys</u> in the AWS Key Management Service Developer Guide.

To use your customer managed key with your Amazon Managed Service for Prometheus workspaces, the following API operations must be permitted in the key policy:

 <u>kms:CreateGrant</u> – Adds a grant to a customer managed key. Grants control access to a specified KMS key, which allows access to grant operations Amazon Managed Service for Prometheus requires. For more information, see <u>Using Grants</u> in the AWS Key Management Service Developer Guide.

This allows Amazon Managed Service for Prometheus to do the following:

- Call GenerateDataKey to generate an encrypted data key and store it, because the data key isn't immediately used to encrypt.
- Call Decrypt to use the stored encrypted data key to access encrypted data.
- <u>kms:DescribeKey</u> Provides the customer managed key details to allow Amazon Managed Service for Prometheus to validate the key.

The following are policy statement examples you can add for Amazon Managed Service for Prometheus:

```
"Statement" : [
   {
     "Sid" : "Allow access to Amazon Managed Service for Prometheus principal within
your account",
     "Effect" : "Allow",
     "Principal" : {
       "AWS" : "*"
     },
     "Action" : [
       "kms:DescribeKey",
       "kms:CreateGrant",
       "kms:GenerateDataKey",
       "kms:Decrypt"
     ],
     "Resource" : "*",
     "Condition" : {
       "StringEquals" : {
         "kms:ViaService" : "aps.region.amazonaws.com",
         "kms:CallerAccount" : "111122223333"
       }
   },
   {
     "Sid": "Allow access for key administrators - not required for Amazon Managed
Service for Prometheus",
     "Effect": "Allow",
     "Principal": {
       "AWS": "arn:aws:iam::111122223333:root"
      },
```

```
"Action" : [
    "kms:*"
    ],
    "Resource": "arn:aws:kms:region:111122223333:key/key_ID"
    },
    <other statements needed for other non-Amazon Managed Service for Prometheus
scenarios>
]
```

- For more information about <u>specifying permissions in a policy</u>, see the AWS Key Management Service Developer Guide.
- For more information about <u>troubleshooting key access</u>, see the AWS Key Management Service Developer Guide.

Step 2: Specifying a customer managed key for Amazon Managed Service for Prometheus

When you create a workspace, you can specify the customer managed key by entering a **KMS Key ARN**, which Amazon Managed Service for Prometheus uses to encrypt the data stored by the workspace.

Step 3: Accessing data from other services, such as Amazon Managed Grafana

This step is optional — it is only required if you need to access your Amazon Managed Service for Prometheus data from another service.

Your encrypted data is not accessible from other services, unless they also have access to use the AWS KMS key. For example, if you want to use Amazon Managed Grafana to create a dashboard or alert on your data, you must give Amazon Managed Grafana access to the key.

To give Amazon Managed Grafana access to your customer managed key

- In your <u>Amazon Managed Grafana workspaces list</u>, select the name for the workspace that you
 want to have access to Amazon Managed Service for Prometheus. This shows you summary
 information about your Amazon Managed Grafana workspace.
- 2. Note the name of the IAM role used by your workspace. The name is in the format AmazonGrafanaServiceRole-<unique-id>. The console shows you the full ARN for the role. You will specify this name in the AWS KMS console in a later step.

- 3. In your <u>AWS KMS Customer managed keys list</u>, choose the customer managed key you used during creation of your Amazon Managed Service for Prometheus workspace. This opens the key configuration details page.
- 4. Next to **Key users**, select the **Add** button.
- 5. From the list of names, choose the Amazon Managed Grafana IAM role that you noted above. To make it easier to find, you can search by the name, as well.
- 6. Choose **Add** to add the IAM role to the list of Key users.

Your Amazon Managed Grafana workspace can now access the data in your Amazon Managed Service for Prometheus workspace. You can add other users or roles to the key users to enable other services to access your workspace.

Amazon Managed Service for Prometheus encryption context

An <u>encryption context</u> is an optional set of key-value pairs that contain additional contextual information about the data.

AWS KMS uses the encryption context as additional authenticated data to support authenticated encryption. When you include an encryption context in a request to encrypt data, AWS KMS binds the encryption context to the encrypted data. To decrypt data, you include the same encryption context in the request.

Amazon Managed Service for Prometheus encryption context

Amazon Managed Service for Prometheus uses the same encryption context in all AWS KMS cryptographic operations, where the key is aws:amp:arn and the value is the <u>Amazon Resource</u> Name (ARN) of the workspace.

Example

```
"encryptionContext": {
    "aws:aps:arn": "arn:aws:aps:us-west-2:111122223333:workspace/ws-sample-1234-
abcd-56ef-7890abcd12ef"
}
```

Using encryption context for monitoring

When you use a symmetric customer managed key to encrypt your workspace data, you can also use the encryption context in audit records and logs to identify how the customer managed key is being used. The encryption context also appears in <u>logs generated by AWS CloudTrail or Amazon</u> <u>CloudWatch Logs</u>.

Using encryption context to control access to your customer managed key

You can use the encryption context in key policies and IAM policies as conditions to control access to your symmetric customer managed key. You can also use encryption context constraints in a grant.

Amazon Managed Service for Prometheus uses an encryption context constraint in grants to control access to the customer managed key in your account or region. The grant constraint requires that the operations that the grant allows use the specified encryption context.

Example

The following are example key policy statements to give access to a customer managed key for a specific encryption context. The condition in this policy statement requires that the grants have an encryption context constraint that specifies the encryption context.

```
{
    "Sid": "Enable DescribeKey",
    "Effect": "Allow",
    "Principal": {
        "AWS": "arn:aws:iam::111122223333:role/ExampleReadOnlyRole"
     },
     "Action": "kms:DescribeKey",
     "Resource": "*"
},
{
     "Sid": "Enable CreateGrant",
     "Effect": "Allow",
     "Principal": {
         "AWS": "arn:aws:iam::111122223333:role/ExampleReadOnlyRole"
     },
     "Action": "kms:CreateGrant",
     "Resource": "*",
     "Condition": {
         "StringEquals": {
             "kms:EncryptionContext:aws:aps:arn": "arn:aws:aps:us-
west-2:111122223333:workspace/ws-sample-1234-abcd-56ef-7890abcd12ef"
          }
     }
```

}

Monitoring your encryption keys for Amazon Managed Service for Prometheus

When you use an AWS KMS customer managed key with your Amazon Managed Service for Prometheus workspaces, you can use <u>AWS CloudTrail</u> or <u>Amazon CloudWatch Logs</u> to track requests that Amazon Managed Service for Prometheus sends to AWS KMS.

The following examples are AWS CloudTrail events for CreateGrant, GenerateDataKey, Decrypt, and DescribeKey to monitor KMS operations called by Amazon Managed Service for Prometheus to access data encrypted by your customer managed key:

CreateGrant

When you use an AWS KMS customer managed key to encrypt your workspace, Amazon Managed Service for Prometheus sends three CreateGrant requests on your behalf to access the KMS key you specified. The grants that Amazon Managed Service for Prometheus creates are specific to the resource associated with the AWS KMS customer managed key.

The following example event records a CreateGrant operation:

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "TESTANDEXAMPLE:Sampleuser01",
        "arn": "arn:aws:sts::111122223333:assumed-role/Admin/Sampleuser01",
        "accountId": "111122223333",
        "accessKeyId": "EXAMPLE-KEY-ID1",
        "sessionContext": {
            "sessionIssuer": {
                "type": "Role",
                "principalId": "TESTANDEXAMPLE:Sampleuser01",
                "arn": "arn:aws:sts::111122223333:assumed-role/Admin/Sampleuser01",
                "accountId": "111122223333",
                "userName": "Admin"
            },
            "webIdFederationData": {},
            "attributes": {
                "mfaAuthenticated": "false",
                "creationDate": "2021-04-22T17:02:00Z"
            }
```

```
User Guide
```

```
},
        "invokedBy": "aps.amazonaws.com"
    },
    "eventTime": "2021-04-22T17:07:02Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "CreateGrant",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "172.12.34.56",
    "userAgent": "ExampleDesktop/1.0 (V1; OS)",
    "requestParameters": {
        "retiringPrincipal": "aps.region.amazonaws.com",
        "operations": [
            "GenerateDataKey",
            "Decrypt",
            "DescribeKey"
        ],
        "keyId": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE",
        "granteePrincipal": "aps.region.amazonaws.com"
    },
    "responseElements": {
        "grantId":
 "0ab0ac0d0b000f00ea00cc0a0e00fc00bce000c000f0000000c0bc0a0000aaafSAMPLE"
    },
    "requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "readOnly": false,
    "resources": [
        {
            "accountId": "111122223333",
            "type": "AWS::KMS::Key",
            "ARN": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "eventCategory": "Management",
    "recipientAccountId": "111122223333"
}
```

GenerateDataKey

When you enable an AWS KMS customer managed key for your workspace, Amazon Managed Service for Prometheus creates a unique key. It sends a GenerateDataKey request to AWS KMS that specifies the AWS KMScustomer managed key for the resource.

The following example event records the GenerateDataKey operation:

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AWSService",
        "invokedBy": "aps.amazonaws.com"
    },
    "eventTime": "2021-04-22T17:07:02Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "GenerateDataKey",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "172.12.34.56",
    "userAgent": "ExampleDesktop/1.0 (V1; OS)",
    "requestParameters": {
        "encryptionContext": {
            "aws:aps:arn": "arn:aws:aps:us-west-2:111122223333:workspace/ws-
sample-1234-abcd-56ef-7890abcd12ef"
        },
        "keySpec": "AES_256",
        "keyId": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE"
    },
    "responseElements": null,
    "requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "readOnly": true,
    "resources": [
        {
            "accountId": "111122223333",
            "type": "AWS::KMS::Key",
            "ARN": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": true,
```

```
"eventCategory": "Management",
    "recipientAccountId": "111122223333",
    "sharedEventID": "57f5dbee-16da-413e-979f-2c4c6663475e"
}
```

Decrypt

When a query is generated on an encrypted workspace, Amazon Managed Service for Prometheus calls the Decrypt operation to use the stored encrypted data key to access the encrypted data.

The following example event records the Decrypt operation:

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AWSService",
        "invokedBy": "aps.amazonaws.com"
    },
    "eventTime": "2021-04-22T17:10:51Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "Decrypt",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "172.12.34.56",
    "userAgent": "ExampleDesktop/1.0 (V1; OS)",
    "requestParameters": {
        "encryptionContext": {
            "aws:aps:arn": "arn:aws:aps:us-west-2:111122223333:workspace/ws-
sample-1234-abcd-56ef-7890abcd12ef"
        },
        "keyId": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE",
        "encryptionAlgorithm": "SYMMETRIC_DEFAULT"
    },
    "responseElements": null,
    "requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "readOnly": true,
    "resources": [
        {
            "accountId": "111122223333",
            "type": "AWS::KMS::Key",
```

```
User Guide
```

DescribeKey

Amazon Managed Service for Prometheus uses the DescribeKey operation to verify if the AWS KMS customer managed key associated with your workspace exists in the account and region.

The following example event records the DescribeKey operation:

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "TESTANDEXAMPLE:Sampleuser01",
        "arn": "arn:aws:sts::111122223333:assumed-role/Admin/Sampleuser01",
        "accountId": "111122223333",
        "accessKeyId": "EXAMPLE-KEY-ID1",
        "sessionContext": {
            "sessionIssuer": {
                "type": "Role",
                "principalId": "TESTANDEXAMPLE:Sampleuser01",
                "arn": "arn:aws:sts::111122223333:assumed-role/Admin/Sampleuser01",
                "accountId": "111122223333",
                "userName": "Admin"
            },
            "webIdFederationData": {},
            "attributes": {
                "mfaAuthenticated": "false",
                "creationDate": "2021-04-22T17:02:00Z"
            }
        },
        "invokedBy": "aps.amazonaws.com"
    },
```

```
"eventTime": "2021-04-22T17:07:02Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "DescribeKey",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "172.12.34.56",
    "userAgent": "ExampleDesktop/1.0 (V1; OS)",
    "requestParameters": {
        "keyId": "00dd0db0-0000-0000-ac00-b0c000SAMPLE"
    },
    "responseElements": null,
    "requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "readOnly": true,
    "resources": [
        {
            "accountId": "111122223333",
            "type": "AWS::KMS::Key",
            "ARN": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "eventCategory": "Management",
    "recipientAccountId": "111122223333"
}
```

Learn more

The following resources provide more information about data encryption at rest.

- For more information about <u>AWS Key Management Service basic concepts</u>, see the AWS Key Management Service Developer Guide.
- For more information about <u>Security best practices for AWS Key Management Service</u>, see the AWS Key Management Service Developer Guide.

Identity and Access Management for Amazon Managed Service for Prometheus

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon Managed Service for Prometheus resources. IAM is an AWS service that you can use with no additional charge.

Topics

- Audience
- Authenticating with identities
- Managing access using policies
- How Amazon Managed Service for Prometheus works with IAM
- Identity-based policy examples for Amazon Managed Service for Prometheus
- Troubleshooting Amazon Managed Service for Prometheus identity and access

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in Amazon Managed Service for Prometheus.

Service user – If you use the Amazon Managed Service for Prometheus service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Amazon Managed Service for Prometheus features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in Amazon Managed Service for Prometheus, see Troubleshooting Amazon Managed Service for Prometheus identity and access.

Service administrator – If you're in charge of Amazon Managed Service for Prometheus resources at your company, you probably have full access to Amazon Managed Service for Prometheus. It's your job to determine which Amazon Managed Service for Prometheus features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Amazon Managed Service for Prometheus, see How Amazon Managed Service for Prometheus works with IAM.

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Amazon Managed Service for Prometheus. To view example Amazon Managed Service for Prometheus identity-based policies that you can use in IAM, see Identity-based policy examples for Amazon Managed Service for Prometheus.

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be *authenticated* (signed in to AWS) as the AWS account root user, as an IAM user, or by assuming an IAM role.

You can sign in to AWS as a federated identity by using credentials provided through an identity source. AWS IAM Identity Center (IAM Identity Center) users, your company's single sign-on authentication, and your Google or Facebook credentials are examples of federated identities. When you sign in as a federated identity, your administrator previously set up identity federation using IAM roles. When you access AWS by using federation, you are indirectly assuming a role.

Depending on the type of user you are, you can sign in to the AWS Management Console or the AWS access portal. For more information about signing in to AWS, see <u>How to sign in to your AWS</u> <u>account</u> in the AWS Sign-In User Guide.

If you access AWS programmatically, AWS provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use AWS tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see <u>AWS Signature Version 4 for API requests</u> in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see <u>Multi-factor authentication</u> in the *AWS IAM Identity Center User Guide* and <u>AWS Multi-factor authentication in IAM</u> in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see <u>Tasks that require root</u> <u>user credentials</u> in the *IAM User Guide*.

Federated identity

As a best practice, require human users, including users that require administrator access, to use federation with an identity provider to access AWS services by using temporary credentials.

A *federated identity* is a user from your enterprise user directory, a web identity provider, the AWS Directory Service, the Identity Center directory, or any user that accesses AWS services by using credentials provided through an identity source. When federated identities access AWS accounts, they assume roles, and the roles provide temporary credentials.

For centralized access management, we recommend that you use AWS IAM Identity Center. You can create users and groups in IAM Identity Center, or you can connect and synchronize to a set of users and groups in your own identity source for use across all your AWS accounts and applications. For information about IAM Identity Center, see <u>What is IAM Identity Center?</u> in the AWS IAM Identity Center User Guide.

IAM users and groups

An <u>IAM user</u> is an identity within your AWS account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see <u>Rotate access keys regularly for use cases that require long-</u> term credentials in the *IAM User Guide*.

An <u>IAM group</u> is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see <u>Use cases for IAM users</u> in the *IAM User Guide*.

IAM roles

An <u>IAM role</u> is an identity within your AWS account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. To temporarily assume an IAM role in the AWS Management Console, you can switch from a user to an IAM role (console). You can assume a

role by calling an AWS CLI or AWS API operation or by using a custom URL. For more information about methods for using roles, see <u>Methods to assume a role</u> in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- Federated user access To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see <u>Create a role for a third-party identity provider</u> (federation) in the *IAM User Guide*. If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the permission set to a role in IAM. For information about permissions sets, see <u>Permission sets</u> in the *AWS IAM Identity Center User Guide*.
- **Temporary IAM user permissions** An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- Cross-account access You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some AWS services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see Cross account resource access in IAM in the IAM User Guide.
- **Cross-service access** Some AWS services use features in other AWS services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
 - Forward access sessions (FAS) When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see Forward access sessions.
 - Service role A service role is an <u>IAM role</u> that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see <u>Create a role to delegate permissions to an AWS service</u> in the *IAM User Guide*.

- Service-linked role A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- Applications running on Amazon EC2 You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see Use an IAM role to grant permissions to applications running on Amazon EC2 instances in the IAM User Guide.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in AWS as JSON documents. For more information about the structure and contents of JSON policy documents, see <u>Overview of JSON policies</u> in the *IAM User Guide*.

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the iam:GetRole action. A user with that policy can get role information from the AWS Management Console, the AWS CLI, or the AWS API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can

perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see Define custom IAM permissions with customer managed policies in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see <u>Choose between managed policies and inline policies</u> in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must <u>specify a principal</u> in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see <u>Access control list (ACL) overview</u> in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

• **Permissions boundaries** – A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user

or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the Principal field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see <u>Permissions boundaries for IAM entities</u> in the *IAM User Guide*.

- Service control policies (SCPs) SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in AWS Organizations. AWS Organizations is a service for grouping and centrally managing multiple AWS accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each AWS account root user. For more information about Organizations and SCPs, see <u>Service</u> <u>control policies</u> in the AWS Organizations User Guide.
- Resource control policies (RCPs) RCPs are JSON policies that you can use to set the maximum available permissions for resources in your accounts without updating the IAM policies attached to each resource that you own. The RCP limits permissions for resources in member accounts and can impact the effective permissions for identities, including the AWS account root user, regardless of whether they belong to your organization. For more information about Organizations and RCPs, including a list of AWS services that support RCPs, see <u>Resource control policies (RCPs)</u> in the AWS Organizations User Guide.
- Session policies Session policies are advanced policies that you pass as a parameter when you
 programmatically create a temporary session for a role or federated user. The resulting session's
 permissions are the intersection of the user or role's identity-based policies and the session
 policies. Permissions can also come from a resource-based policy. An explicit deny in any of these
 policies overrides the allow. For more information, see <u>Session policies</u> in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see <u>Policy evaluation logic</u> in the *IAM User Guide*.

How Amazon Managed Service for Prometheus works with IAM

Before you use IAM to manage access to Amazon Managed Service for Prometheus, learn what IAM features are available to use with Amazon Managed Service for Prometheus.

IAM features you can use with Amazon Managed Service for Prometheus

IAM feature	Amazon Managed Service for Prometheus support
Identity-based policies	Yes
Resource-based policies	No
Policy actions	Yes
Policy resources	Yes
Policy condition keys	No
ACLs	No
ABAC (tags in policies)	Yes
Temporary credentials	Yes
Forward access sessions (FAS)	No
Service roles	No
Service-linked roles	Yes

To get a high-level view of how Amazon Managed Service for Prometheus and other AWS services work with most IAM features, see <u>AWS services that work with IAM</u> in the *IAM User Guide*.

Identity-based policies for Amazon Managed Service for Prometheus

Supports identity-based policies: Yes

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see Define custom IAM permissions with customer managed policies in the *IAM User Guide*.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. You can't specify the principal in an

identity-based policy because it applies to the user or role to which it is attached. To learn about all of the elements that you can use in a JSON policy, see <u>IAM JSON policy elements reference</u> in the *IAM User Guide*.

Identity-based policy examples for Amazon Managed Service for Prometheus

To view examples of Amazon Managed Service for Prometheus identity-based policies, see Identity-based policy examples for Amazon Managed Service for Prometheus.

Resource-based policies within Amazon Managed Service for Prometheus

Supports resource-based policies: No

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must <u>specify a principal</u> in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. Adding a cross-account principal to a resource-based policy is only half of establishing the trust relationship. When the principal and the resource are in different AWS accounts, an IAM administrator in the trusted account must also grant the principal entity (user or role) permission to access the resource. They grant permission by attaching an identity-based policy to the entity. However, if a resource-based policy grants access to a principal in the same account, no additional identity-based policy is required. For more information, see <u>Cross account resource access in IAM</u> in the *IAM User Guide*.

Policy actions for Amazon Managed Service for Prometheus

Supports policy actions: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Action element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated AWS API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API

operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

To see a list of Amazon Managed Service for Prometheus actions, see <u>Actions defined by Amazon</u> <u>Managed Service for Prometheus</u> in the *Service Authorization Reference*.

Policy actions in Amazon Managed Service for Prometheus use the following prefix before the action:

aps

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [
"aps:action1",
"aps:action2"
]
```

To view examples of Amazon Managed Service for Prometheus identity-based policies, see Identity-based policy examples for Amazon Managed Service for Prometheus.

Policy resources for Amazon Managed Service for Prometheus

Supports policy resources: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its <u>Amazon Resource Name (ARN)</u>. You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

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

To see a list of Amazon Managed Service for Prometheus resource types and their ARNs, see <u>Resources defined by Amazon Managed Service for Prometheus</u> in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see <u>Actions</u> defined by Amazon Managed Service for Prometheus.

To view examples of Amazon Managed Service for Prometheus identity-based policies, see Identity-based policy examples for Amazon Managed Service for Prometheus.

Policy condition keys for Amazon Managed Service for Prometheus

Supports service-specific policy condition keys: No

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Condition element (or Condition *block*) lets you specify conditions in which a statement is in effect. The Condition element is optional. You can create conditional expressions that use <u>condition operators</u>, such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple Condition elements in a statement, or multiple keys in a single Condition element, AWS evaluates them using a logical AND operation. If you specify multiple values for a single condition key, AWS evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see IAM policy elements: variables and tags in the *IAM User Guide*.

AWS supports global condition keys and service-specific condition keys. To see all AWS global condition keys, see <u>AWS global condition context keys</u> in the *IAM User Guide*.

To see a list of Amazon Managed Service for Prometheus condition keys, see <u>Condition keys for</u> <u>Amazon Managed Service for Prometheus</u> in the *Service Authorization Reference*. To learn with which actions and resources you can use a condition key, see <u>Actions defined by Amazon Managed</u> <u>Service for Prometheus</u>.

To view examples of Amazon Managed Service for Prometheus identity-based policies, see Identity-based policy examples for Amazon Managed Service for Prometheus.

Access control lists (ACLs) in Amazon Managed Service for Prometheus

Supports ACLs: No

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Attribute-based access control (ABAC) with Amazon Managed Service for Prometheus

Supports ABAC (tags in policies): Yes

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. In AWS, these attributes are called *tags*. You can attach tags to IAM entities (users or roles) and to many AWS resources. 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 are trying to access.

ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

To control access based on tags, you provide 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.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see <u>Define permissions with ABAC authorization</u> in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see <u>Use attribute-based access control</u> (ABAC) in the *IAM User Guide*.

Using temporary credentials with Amazon Managed Service for Prometheus

Supports temporary credentials: Yes

Some AWS services don't work when you sign in using temporary credentials. For additional information, including which AWS services work with temporary credentials, see <u>AWS services that</u> work with IAM in the *IAM User Guide*.

You are using temporary credentials if you sign in to the AWS Management Console using any method except a user name and password. For example, when you access AWS using your company's single sign-on (SSO) link, that process automatically creates temporary credentials. You also automatically create temporary credentials when you sign in to the console as a user and then switch roles. For more information about switching roles, see <u>Switch from a user to an IAM role</u> (console) in the *IAM User Guide*.

You can manually create temporary credentials using the AWS CLI or AWS API. You can then use those temporary credentials to access AWS. AWS recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see <u>Temporary security credentials in IAM</u>.

Forward access sessions for Amazon Managed Service for Prometheus

Supports forward access sessions (FAS): No

When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see Forward access sessions.

Service roles for Amazon Managed Service for Prometheus

Supports service roles: No

A service role is an <u>IAM role</u> that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see <u>Create a role to delegate permissions to an AWS service</u> in the *IAM User Guide*.

<u> Marning</u>

Changing the permissions for a service role might break Amazon Managed Service for Prometheus functionality. Edit service roles only when Amazon Managed Service for Prometheus provides guidance to do so.

Service-linked roles for Amazon Managed Service for Prometheus

Supports service-linked roles: Yes

A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For details about creating or managing Amazon Managed Service for Prometheus service-linked roles, see Using service-linked roles for Amazon Managed Service for Prometheus.

Identity-based policy examples for Amazon Managed Service for Prometheus

By default, users and roles don't have permission to create or modify Amazon Managed Service for Prometheus resources. They also can't perform tasks by using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS API. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see <u>Create IAM policies (console)</u> in the *IAM User Guide*.

For details about actions and resource types defined by Amazon Managed Service for Prometheus, including the format of the ARNs for each of the resource types, see <u>Actions, resources, and</u> <u>condition keys for Amazon Managed Service for Prometheus</u> in the *Service Authorization Reference*.

Topics

- Policy best practices
- Using the Amazon Managed Service for Prometheus console
- Allow users to view their own permissions

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Amazon Managed Service for Prometheus resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- Get started with AWS managed policies and move toward least-privilege permissions To get started granting permissions to your users and workloads, use the AWS managed policies that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see <u>AWS managed policies</u> or <u>AWS</u> <u>managed policies for job functions</u> in the *IAM User Guide*.
- **Apply least-privilege permissions** When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see <u>Policies and permissions in IAM</u> in the *IAM User Guide*.
- Use conditions in IAM policies to further restrict access You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific AWS service, such as AWS CloudFormation. For more information, see <u>IAM JSON policy elements: Condition</u> in the *IAM User Guide*.
- Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see <u>Validate policies with IAM Access Analyzer</u> in the *IAM User Guide*.
- Require multi-factor authentication (MFA) If you have a scenario that requires IAM users or a root user in your AWS account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see <u>Secure API</u> access with MFA in the IAM User Guide.

For more information about best practices in IAM, see <u>Security best practices in IAM</u> in the *IAM User Guide*.

Using the Amazon Managed Service for Prometheus console

To access the Amazon Managed Service for Prometheus console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Amazon Managed Service for Prometheus resources in your AWS account. If you create an identity-based policy that

is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

To ensure that users and roles can still use the Amazon Managed Service for Prometheus console, also attach the Amazon Managed Service for Prometheus ConsoleAccess or ReadOnly AWS managed policy to the entities. For more information, see <u>Adding permissions to a user</u> in the *IAM User Guide*.

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ViewOwnUserInfo",
            "Effect": "Allow",
            "Action": [
                "iam:GetUserPolicy",
                "iam:ListGroupsForUser",
                "iam:ListAttachedUserPolicies",
                "iam:ListUserPolicies",
                "iam:GetUser"
            ],
            "Resource": ["arn:aws:iam::*:user/${aws:username}"]
        },
        {
            "Sid": "NavigateInConsole",
            "Effect": "Allow",
            "Action": [
                "iam:GetGroupPolicy",
                "iam:GetPolicyVersion",
                "iam:GetPolicy",
                "iam:ListAttachedGroupPolicies",
                "iam:ListGroupPolicies",
```

```
"iam:ListPolicyVersions",
    "iam:ListPolicies",
    "iam:ListUsers"
    ],
    "Resource": "*"
    }
]
```

Troubleshooting Amazon Managed Service for Prometheus identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon Managed Service for Prometheus and IAM.

Topics

- I am not authorized to perform an action in Amazon Managed Service for Prometheus
- I am not authorized to perform iam:PassRole
- I want to allow people outside of my AWS account to access my Amazon Managed Service for Prometheus resources

I am not authorized to perform an action in Amazon Managed Service for Prometheus

If you receive an error that you're not authorized to perform an action, your policies must be updated to allow you to perform the action.

The following example error occurs when the mateojackson IAM user tries to use the console to view details about a fictional *my*-*example*-*widget* resource but doesn't have the fictional aps: *GetWidget* permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
    aps:GetWidget on resource: my-example-widget
```

In this case, the policy for the mateojackson user must be updated to allow access to the *myexample-widget* resource by using the aps: *GetWidget* action. If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the iam: PassRole action, your policies must be updated to allow you to pass a role to Amazon Managed Service for Prometheus.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named marymajor tries to use the console to perform an action in Amazon Managed Service for Prometheus. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform:
  iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the iam: PassRole action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my AWS account to access my Amazon Managed Service for Prometheus resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Amazon Managed Service for Prometheus supports these features, see <u>How</u> Amazon Managed Service for Prometheus works with IAM.
- To learn how to provide access to your resources across AWS accounts that you own, see Providing access to an IAM user in another AWS account that you own in the *IAM User Guide*.

- To learn how to provide access to your resources to third-party AWS accounts, see <u>Providing</u> access to AWS accounts owned by third parties in the *IAM User Guide*.
- To learn how to provide access through identity federation, see <u>Providing access to externally</u> authenticated users (identity federation) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see Cross account resource access in IAM in the *IAM User Guide*.

IAM permissions and policies

Access to Amazon Managed Service for Prometheus actions and data requires credentials. Those credentials must have permissions to perform the actions and to access the AWS resources, such as retrieving Amazon Managed Service for Prometheus data about your cloud resources. The following sections provide details about how you can use AWS Identity and Access Management (IAM) and Amazon Managed Service for Prometheus to help secure your resources, by controlling who can access them. For more information, see Policies and permissions in IAM.

Amazon Managed Service for Prometheus permissions

To see the list of possible Amazon Managed Service for Prometheus actions. resource types, and condition keys, see <u>Actions, resources, and condition keys for Amazon Managed Service for</u> <u>Prometheus</u>.

Sample IAM policies

This section provides examples of other self-managed policies that you can create.

The following IAM policy grants full access to Amazon Managed Service for Prometheus and also enables a user to discover Amazon EKS clusters and see the details about them.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
               "aps:*",
               "eks:DescribeCluster",
               "eks:ListClusters"
        ],
```

```
"Resource": "*"
}
]
```

}

Compliance Validation for Amazon Managed Service for Prometheus

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 AWS Compliance Programs.

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 Compliance & Governance</u> These solution implementation guides discuss architectural considerations and provide steps for deploying security and compliance features.
- <u>HIPAA Eligible Services Reference</u> Lists HIPAA eligible services. Not all AWS services are HIPAA eligible.
- <u>AWS Compliance Resources</u> This collection of workbooks and guides might apply to your industry and location.
- <u>AWS Customer Compliance Guides</u> 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)).
- <u>Evaluating Resources with Rules</u> in the *AWS Config Developer Guide* The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- <u>AWS Security Hub</u> 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.

- <u>Amazon GuardDuty</u> 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.
- <u>AWS Audit Manager</u> 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 Amazon Managed Service for Prometheus

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.

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

In addition to the AWS global infrastructure, Amazon Managed Service for Prometheus offers several features to help support your data resiliency and backup needs, including support for <u>high-availability data</u>.

Infrastructure Security in Amazon Managed Service for Prometheus

As a managed service, Amazon Managed Service for Prometheus 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 Protection</u> in *Security Pillar AWS Well-Architected Framework*.

You use AWS published API calls to access Amazon Managed Service for Prometheus 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.

Using service-linked roles for Amazon Managed Service for Prometheus

Amazon Managed Service for Prometheus uses AWS Identity and Access Management (IAM) <u>service-linked roles</u>. A service-linked role is a unique type of IAM role that is linked directly to Amazon Managed Service for Prometheus. Service-linked roles are predefined by Amazon Managed Service for Prometheus and include all the permissions that the service requires to call other AWS services on your behalf.

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

Using roles for scraping metrics from EKS

When automatically scraping metrics using Amazon Managed Service for Prometheus managed collector, the AWSServiceRoleForAmazonPrometheusScraper service-linked role is used to make setting up managed collector easier, because you don't have to manually add the necessary permissions. Amazon Managed Service for Prometheus defines the permissions, and only Amazon Managed Service for Prometheus the role.

For information about other services that support service-linked roles, see <u>AWS services that work</u> <u>with IAM</u> and look for the services that have **Yes** in the **Service-linked roles** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

Service-linked role permissions for Amazon Managed Service for Prometheus

Amazon Managed Service for Prometheus uses a service-linked role named with the prefix **AWSServiceRoleForAmazonPrometheusScraper** to allow Amazon Managed Service for Prometheus to automatically scrape metrics in your Amazon EKS clusters.

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

scraper.aps.amazonaws.com

The role permissions policy named AmazonPrometheusScraperServiceRolePolicy allows Amazon Managed Service for Prometheus to complete the following actions on the specified resources:

- Ready and modify network configuration to connect to the network that contains your Amazon EKS cluster.
- Read metrics from Amazon EKS clusters and write metrics to your Amazon Managed Service for Prometheus workspaces.

You must configure permissions to allow your users, groups, or roles to create a service-linked role. For more information, see <u>Service-linked role permissions</u> in the *IAM User Guide*.

Creating a service-linked role for Amazon Managed Service for Prometheus

You don't need to manually create a service-linked role. When you create an managed collector instance using Amazon EKS or Amazon Managed Service for Prometheus in the AWS Management Console, the AWS CLI, or the AWS API, Amazon Managed Service for Prometheus creates the service-linked role for you.

🔥 Important

This service-linked role can appear in your account if you completed an action in another service that uses the features supported by this role. To learn more, see <u>A new role</u> appeared in my AWS account.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you create an managed collector instance using Amazon

EKS or Amazon Managed Service for Prometheus, Amazon Managed Service for Prometheus creates the service-linked role for you again.

Editing a service-linked role for Amazon Managed Service for Prometheus

Amazon Managed Service for Prometheus does not allow you to edit the AWSServiceRoleForAmazonPrometheusScraper service-linked role. After you create a servicelinked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see <u>Editing</u> <u>a service-linked role</u> in the *IAM User Guide*.

Deleting a service-linked role for Amazon Managed Service for Prometheus

You don't need to manually delete the AWSServiceRoleForAmazonPrometheusScraper role. When you delete all managed collector instances associated with the role in the AWS Management Console, the AWS CLI, or the AWS API, Amazon Managed Service for Prometheus cleans up the resources and deletes the service-linked role for you.

Supported Regions for Amazon Managed Service for Prometheus service-linked roles

Amazon Managed Service for Prometheus supports using service-linked roles in all of the Regions where the service is available. For more information, see <u>Supported Regions</u>.

Logging Amazon Managed Service for Prometheus API calls using AWS CloudTrail

Amazon Managed Service for Prometheus is integrated with <u>AWS CloudTrail</u>, a service that provides a record of actions taken by a user, role, or an AWS service. CloudTrail captures all API calls for Amazon Managed Service for Prometheus as events. The calls captured include calls from the Amazon Managed Service for Prometheus console and code calls to the Amazon Managed Service for Prometheus API operations. Using the information collected by CloudTrail, you can determine the request that was made to Amazon Managed Service for Prometheus, the IP address from which the request was made, when it was made, and additional details.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root user or user credentials.
- Whether the request was made on behalf of an IAM Identity Center user.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

CloudTrail is active in your AWS account when you create the account and you automatically have access to the CloudTrail **Event history**. The CloudTrail **Event history** provides a viewable, searchable, downloadable, and immutable record of the past 90 days of recorded management events in an AWS Region. For more information, see <u>Working with CloudTrail Event history</u> in the *AWS CloudTrail User Guide*. There are no CloudTrail charges for viewing the **Event history**.

For an ongoing record of events in your AWS account past 90 days, create a trail or a <u>CloudTrail</u> <u>Lake</u> event data store.

CloudTrail trails

A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. All trails created using the AWS Management Console are multi-Region. You can create a single-Region or a multi-Region trail by using the AWS CLI. Creating a multi-Region trail is recommended because you capture activity in all AWS Regions in your account. If you create a single-Region trail, you can view only the events logged in the trail's AWS Region. For more information about trails, see <u>Creating a trail for your AWS account</u> and <u>Creating a trail for an organization</u> in the AWS CloudTrail User *Guide*.

You can deliver one copy of your ongoing management events to your Amazon S3 bucket at no charge from CloudTrail by creating a trail, however, there are Amazon S3 storage charges. For more information about CloudTrail pricing, see <u>AWS CloudTrail Pricing</u>. For information about Amazon S3 pricing, see <u>Amazon S3 Pricing</u>.

CloudTrail Lake event data stores

CloudTrail Lake lets you run SQL-based queries on your events. CloudTrail Lake converts existing events in row-based JSON format to <u>Apache ORC</u> format. ORC is a columnar storage format that is optimized for fast retrieval of data. Events are aggregated into *event data stores*, which are immutable collections of events based on criteria that you select by applying <u>advanced</u> <u>event selectors</u>. The selectors that you apply to an event data store control which events persist and are available for you to query. For more information about CloudTrail Lake, see <u>Working</u> with AWS CloudTrail Lake in the AWS CloudTrail User Guide.

CloudTrail Lake event data stores and queries incur costs. When you create an event data store, you choose the <u>pricing option</u> you want to use for the event data store. The pricing option determines the cost for ingesting and storing events, and the default and maximum retention period for the event data store. For more information about CloudTrail pricing, see <u>AWS CloudTrail Pricing</u>.

Amazon Managed Service for Prometheus management events in CloudTrail

<u>Management events</u> provide information about management operations that are performed on resources in your AWS account. These are also known as control plane operations. By default, CloudTrail logs management events.

Amazon Managed Service for Prometheus logs all Amazon Managed Service for Prometheus control plane operations as management events. For a list of the Amazon Managed Service for Prometheus control plane operations that Amazon Managed Service for Prometheus logs to CloudTrail, see the Amazon Managed Service for Prometheus API Reference.

Amazon Managed Service for Prometheus event examples

An event represents a single request from any source and includes information about the requested API operation, the date and time of the operation, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so events don't appear in any specific order.

Example: CreateWorkspace

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

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "EXAMPLE123EXAMPLE123-1234567890616",
        "arn": "arn:aws:sts::123456789012:assumed-role/Admin/admin",
        "accountId": "123456789012",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "sessionContext": {
            "sessionIssuer": {
            "sessionIssuer}
            "sessionIssuer": {
            "s
```

```
"type": "Role",
                "principalId": "AIDACKCEVSQ6C2EXAMPLE",
                "arn": "arn:aws:iam::123456789012:role/Admin",
                "accountId": "123456789012",
                "userName": "Admin"
            },
            "webIdFederationData": {
            },
            "attributes": {
                "mfaAuthenticated": "false",
                "creationDate": "2020-11-30T23:39:29Z"
            }
        }
    },
    "eventTime": "2020-11-30T23:43:21Z",
    "eventSource": "aps.amazonaws.com",
    "eventName": "CreateWorkspace",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "203.0.113.1",
    "userAgent": "aws-cli/1.11.167 Python/2.7.10 Darwin/16.7.0 botocore/1.7.25",
    "requestParameters": {
        "alias": "alias-example",
        "clientToken": "12345678-1234-abcd-1234-12345abcd1"
    },
    "responseElements": {
        "Access-Control-Expose-Headers": "x-amzn-errortype,x-amzn-requestid,x-amzn-
trace-id, x-amzn-errormessage, x-amz-apigw-id, date",
        "arn": "arn:aws:aps:us-west-2:123456789012:workspace/ws-abc123456-
abcd-1234-5678-1234567890",
        "status": {
            "statusCode": "CREATING"
        },
        "workspaceId": "ws-12345678-1234-abcd-1234-1234567890"
    },
    "requestID": "890b8639-e51f-11e7-b038-EXAMPLE",
    "eventID": "874f89fa-70fc-4798-bc00-EXAMPLE",
    "readOnly": false,
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "eventCategory": "Management",
    "recipientAccountId": "123456789012"
}
```

Example: CreateAlertManagerDefinition

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

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "EXAMPLE123EXAMPLE123-1234567890616",
        "arn": "arn:aws:sts::123456789012:assumed-role/Admin/admin",
        "accountId": "123456789012",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "sessionContext": {
            "sessionIssuer": {
                "type": "Role",
                "principalId": "AIDACKCEVSQ6C2EXAMPLE",
                "arn": "arn:aws:iam::123456789012:role/Admin",
                "accountId": "123456789012",
                "userName": "Admin"
            },
            "webIdFederationData": {
            },
            "attributes": {
                "mfaAuthenticated": "false",
                "creationDate": "2021-09-23T20:20:14Z"
            }
        }
    },
    "eventTime": "2021-09-23T20:22:43Z",
    "eventSource": "aps.amazonaws.com",
    "eventName": "CreateAlertManagerDefinition",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "203.0.113.1",
    "userAgent": "Boto3/1.17.46 Python/3.6.14 Linux/4.14.238-182.422.amzn2.x86_64 exec-
env/AWS_ECS_FARGATE Botocore/1.20.46",
    "requestParameters": {
        "data":
 "YWxlcnRtYW5hZ2VyX2NvbmZpZzoqfAoqIGdsb2JhbDoKICAqIHNtdHBfc21hcnRob3N00iAnbG9jYWxob3N00jI1JwoqI
        "clientToken": "12345678-1234-abcd-1234-12345abcd1",
        "workspaceId": "ws-12345678-1234-abcd-1234-1234567890"
    },
    "responseElements": {
```

```
"Access-Control-Expose-Headers": "x-amzn-errortype,x-amzn-requestid,x-amzn-
trace-id,x-amzn-errormessage,x-amz-apigw-id,date",
    "status": {
        "statusCode": "CREATING"
      }
    },
    "requestID": "890b8639-e51f-11e7-b038-EXAMPLE",
    "requestID": "874f89fa-70fc-4798-bc00-EXAMPLE",
    "readOnly": false,
    "eventID": "AwsApiCall",
    "managementEvent": true,
    "eventType": "AwsApiCall",
    "menagementEvent": true,
    "eventCategory": "Management",
    "recipientAccountId": "123456789012"
}
```

Example: CreateRuleGroupsNamespace

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

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "EXAMPLE123EXAMPLE123-1234567890616",
        "arn": "arn:aws:sts::123456789012:assumed-role/Admin/admin",
        "accountId": "123456789012",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "sessionContext": {
            "sessionIssuer": {
                "type": "Role",
                "principalId": "AIDACKCEVSQ6C2EXAMPLE",
                "arn": "arn:aws:iam::123456789012:role/Admin",
                "accountId": "123456789012",
                "userName": "Admin"
            },
            "webIdFederationData": {
            },
            "attributes": {
                "creationDate": "2021-09-23T20:22:19Z",
                "mfaAuthenticated": "false"
            }
```

```
}
    },
    "eventTime": "2021-09-23T20:25:08Z",
    "eventSource": "aps.amazonaws.com",
    "eventName": "CreateRuleGroupsNamespace",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "34.212.33.165",
    "userAgent": "Boto3/1.17.63 Python/3.6.14 Linux/4.14.238-182.422.amzn2.x86_64 exec-
env/AWS_ECS_FARGATE Botocore/1.20.63",
    "requestParameters": {
        "data":
 "Z3JvdXBz0gogIC0gbmFtZTogdGVzdFJ1bGVHcm91cHN0YW11c3BhY2UKICAgIHJ1bGVz0gogICAgLSBhbGVydDogdGVzc
        "clientToken": "12345678-1234-abcd-1234-12345abcd1",
        "name": "exampleRuleGroupsNamespace",
        "workspaceId": "ws-12345678-1234-abcd-1234-1234567890"
    },
    "responseElements": {
        "Access-Control-Expose-Headers": "x-amzn-errortype,x-amzn-requestid,x-amzn-
trace-id, x-amzn-errormessage, x-amz-apigw-id, date",
        "name": "exampleRuleGroupsNamespace",
        "arn": "arn:aws:aps:us-west-2:492980759322:rulegroupsnamespace/ws-
ae46a85c-1609-4c22-90a3-2148642c3b6c/exampleRuleGroupsNamespace",
        "status": {
            "statusCode": "CREATING"
        },
        "tags": {}
    },
    "requestID": "890b8639-e51f-11e7-b038-EXAMPLE",
    "eventID": "874f89fa-70fc-4798-bc00-EXAMPLE",
    "readOnly": false,
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "eventCategory": "Management",
    "recipientAccountId": "123456789012"
}
```

For information about CloudTrail record contents, see <u>CloudTrail record contents</u> in the AWS *CloudTrail User Guide*.

Set up IAM roles for service accounts

With IAM roles for service accounts, you can associate an IAM role with a Kubernetes service account. This service account can then provide AWS permissions to the containers in any pod that uses that service account. For more information, see IAM roles for service accounts.

IAM roles for service accounts are also known as service roles.

In Amazon Managed Service for Prometheus, using service roles can help you get the roles you need to authorize and authenticate between Amazon Managed Service for Prometheus, Prometheus servers, and Grafana servers.

Prerequisites

The procedures on this page require that you have the AWS CLI and EKSCTL command line interface installed.

Set up service roles for the ingestion of metrics from Amazon EKS clusters

To set up the service roles to enable Amazon Managed Service for Prometheus to ingest metrics from Prometheus servers in Amazon EKS clusters, you must be logged on to an account with the following permissions:

- iam:CreateRole
- iam:CreatePolicy
- iam:GetRole
- iam:AttachRolePolicy
- iam:GetOpenIDConnectProvider

To set up the service role for ingestion into Amazon Managed Service for Prometheus

 Create a file named createIRSA-AMPIngest.sh with the following content. Replace <my_amazon_eks_clustername> with the name of your cluster, and replace <my_prometheus_namespace> with your Prometheus namespace.

```
#!/bin/bash -e
CLUSTER_NAME=<my_amazon_eks_clustername>
```

```
SERVICE_ACCOUNT_NAMESPACE=<my_prometheus_namespace>
AWS_ACCOUNT_ID=$(aws sts get-caller-identity --query "Account" --output text)
OIDC_PROVIDER=$(aws eks describe-cluster --name $CLUSTER_NAME --query
 "cluster.identity.oidc.issuer" --output text | sed -e "s/^https:\/\//")
SERVICE_ACCOUNT_AMP_INGEST_NAME=amp-iamproxy-ingest-service-account
SERVICE_ACCOUNT_IAM_AMP_INGEST_ROLE=amp-iamproxy-ingest-role
SERVICE_ACCOUNT_IAM_AMP_INGEST_POLICY=AMPIngestPolicy
#
# Set up a trust policy designed for a specific combination of K8s service account
and namespace to sign in from a Kubernetes cluster which hosts the OIDC Idp.
#
cat <<EOF > TrustPolicy.json
Ł
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Federated": "arn:aws:iam::${AWS_ACCOUNT_ID}:oidc-provider/
${OIDC_PROVIDER}"
      },
      "Action": "sts:AssumeRoleWithWebIdentity",
      "Condition": {
        "StringEquals": {
          "${OIDC_PROVIDER}:sub": "system:serviceaccount:
${SERVICE_ACCOUNT_NAMESPACE}:${SERVICE_ACCOUNT_AMP_INGEST_NAME}"
        }
      }
    }
  ]
}
EOF
#
# Set up the permission policy that grants ingest (remote write) permissions for
all AMP workspaces
#
cat <<EOF > PermissionPolicyIngest.json
{
  "Version": "2012-10-17",
   "Statement": [
       {"Effect": "Allow",
        "Action": [
           "aps:RemoteWrite",
           "aps:GetSeries",
```

```
"aps:GetLabels",
           "aps:GetMetricMetadata"
        ],
        "Resource": "*"
      }
  ]
}
EOF
function getRoleArn() {
  OUTPUT=$(aws iam get-role --role-name $1 --query 'Role.Arn' --output text 2>&1)
  # Check for an expected exception
  if [[ $? -eq 0 ]]; then
    echo $0UTPUT
  elif [[ -n $(grep "NoSuchEntity" <<< $OUTPUT) ]]; then</pre>
   echo ""
  else
   >&2 echo $OUTPUT
   return 1
 fi
}
#
# Create the IAM Role for ingest with the above trust policy
#
SERVICE_ACCOUNT_IAM_AMP_INGEST_ROLE_ARN=$(getRoleArn
$SERVICE_ACCOUNT_IAM_AMP_INGEST_ROLE)
if [ "$SERVICE_ACCOUNT_IAM_AMP_INGEST_ROLE_ARN" = "" ];
then
  #
  # Create the IAM role for service account
  #
  SERVICE_ACCOUNT_IAM_AMP_INGEST_ROLE_ARN=$(aws iam create-role \
  --role-name $SERVICE_ACCOUNT_IAM_AMP_INGEST_ROLE \
  --assume-role-policy-document file://TrustPolicy.json \
  --query "Role.Arn" --output text)
  #
  # Create an IAM permission policy
  #
  SERVICE_ACCOUNT_IAM_AMP_INGEST_ARN=$(aws iam create-policy --policy-name
 $SERVICE_ACCOUNT_IAM_AMP_INGEST_POLICY \
  --policy-document file://PermissionPolicyIngest.json \
  --query 'Policy.Arn' --output text)
```

#

```
# Attach the required IAM policies to the IAM role created above
 #
  aws iam attach-role-policy ∖
  --role-name $SERVICE_ACCOUNT_IAM_AMP_INGEST_ROLE \
  --policy-arn $SERVICE_ACCOUNT_IAM_AMP_INGEST_ARN
else
    echo "$SERVICE_ACCOUNT_IAM_AMP_INGEST_ROLE_ARN IAM role for ingest already
exists"
fi
echo $SERVICE_ACCOUNT_IAM_AMP_INGEST_ROLE_ARN
#
# EKS cluster hosts an OIDC provider with a public discovery endpoint.
# Associate this IdP with AWS IAM so that the latter can validate and accept the
OIDC tokens issued by Kubernetes to service accounts.
# Doing this with eksctl is the easier and best approach.
#
eksctl utils associate-iam-oidc-provider --cluster $CLUSTER_NAME --approve
```

2. Enter the following command to give the script the necessary privileges.

```
chmod +x createIRSA-AMPIngest.sh
```

3. Run the script.

Set up IAM roles for service accounts for the querying of metrics

To set up the IAM role for service account (service role) to enable the querying of metrics from Amazon Managed Service for Prometheus workspaces, you must be logged on to an account with the following permissions:

- iam:CreateRole
- iam:CreatePolicy
- iam:GetRole
- iam:AttachRolePolicy
- iam:GetOpenIDConnectProvider

To set up service roles for the querying of Amazon Managed Service for Prometheus metrics;

 Create a file named createIRSA-AMPQuery.sh with the following content. Replace <my_amazon_eks_clustername> with the name of your cluster, and replace
 <my_prometheus_namespace> with your Prometheus namespace.

```
#!/bin/bash -e
CLUSTER_NAME=<my_amazon_eks_clustername>
SERVICE_ACCOUNT_NAMESPACE=<my_prometheus_namespace>
AWS_ACCOUNT_ID=$(aws sts get-caller-identity --query "Account" --output text)
OIDC_PROVIDER=$(aws eks describe-cluster --name $CLUSTER_NAME --query
 "cluster.identity.oidc.issuer" --output text | sed -e "s/^https:\/\///")
SERVICE_ACCOUNT_AMP_QUERY_NAME=amp-iamproxy-query-service-account
SERVICE_ACCOUNT_IAM_AMP_QUERY_ROLE=amp-iamproxy-query-role
SERVICE_ACCOUNT_IAM_AMP_QUERY_POLICY=AMPQueryPolicy
#
# Setup a trust policy designed for a specific combination of K8s service account
and namespace to sign in from a Kubernetes cluster which hosts the OIDC Idp.
#
cat <<EOF > TrustPolicy.json
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Federated": "arn:aws:iam::${AWS_ACCOUNT_ID}:oidc-provider/
${OIDC_PROVIDER}"
      },
      "Action": "sts:AssumeRoleWithWebIdentity",
      "Condition": {
        "StringEquals": {
          "${OIDC_PROVIDER}:sub": "system:serviceaccount:
${SERVICE_ACCOUNT_NAMESPACE}:${SERVICE_ACCOUNT_AMP_QUERY_NAME}"
        }
      }
    }
  ]
}
EOF
#
# Set up the permission policy that grants query permissions for all AMP workspaces
#
```

```
User Guide
```

```
cat <<EOF > PermissionPolicyQuery.json
{
  "Version": "2012-10-17",
   "Statement": [
       {"Effect": "Allow",
        "Action": [
           "aps:QueryMetrics",
           "aps:GetSeries",
           "aps:GetLabels",
           "aps:GetMetricMetadata"
        ],
        "Resource": "*"
      }
   ]
}
EOF
function getRoleArn() {
  OUTPUT=$(aws iam get-role --role-name $1 --query 'Role.Arn' --output text 2>&1)
  # Check for an expected exception
  if [[ $? -eq 0 ]]; then
    echo $0UTPUT
  elif [[ -n $(grep "NoSuchEntity" <<< $OUTPUT) ]]; then</pre>
    echo ""
  else
    >&2 echo $OUTPUT
    return 1
 fi
}
#
# Create the IAM Role for query with the above trust policy
#
SERVICE_ACCOUNT_IAM_AMP_QUERY_ROLE_ARN=$(getRoleArn
$SERVICE_ACCOUNT_IAM_AMP_QUERY_ROLE)
if [ "$SERVICE_ACCOUNT_IAM_AMP_QUERY_ROLE_ARN" = "" ];
then
  #
  # Create the IAM role for service account
  #
  SERVICE_ACCOUNT_IAM_AMP_QUERY_ROLE_ARN=$(aws iam create-role \
  --role-name $SERVICE_ACCOUNT_IAM_AMP_QUERY_ROLE \
```

```
--assume-role-policy-document file://TrustPolicy.json \setminus
```

```
--query "Role.Arn" --output text)
 #
 # Create an IAM permission policy
 #
 SERVICE_ACCOUNT_IAM_AMP_QUERY_ARN=$(aws iam create-policy --policy-name
 $SERVICE_ACCOUNT_IAM_AMP_QUERY_POLICY \
  --policy-document file://PermissionPolicyQuery.json \
  --query 'Policy.Arn' --output text)
 #
 # Attach the required IAM policies to the IAM role create above
 #
 aws iam attach-role-policy \setminus
  --role-name $SERVICE_ACCOUNT_IAM_AMP_QUERY_ROLE \
  --policy-arn $SERVICE_ACCOUNT_IAM_AMP_QUERY_ARN
else
    echo "$SERVICE_ACCOUNT_IAM_AMP_QUERY_ROLE_ARN IAM role for query already
exists"
fi
echo $SERVICE_ACCOUNT_IAM_AMP_QUERY_ROLE_ARN
#
# EKS cluster hosts an OIDC provider with a public discovery endpoint.
# Associate this IdP with AWS IAM so that the latter can validate and accept the
OIDC tokens issued by Kubernetes to service accounts.
# Doing this with eksctl is the easier and best approach.
#
eksctl utils associate-iam-oidc-provider --cluster $CLUSTER_NAME --approve
```

2. Enter the following command to give the script the necessary privileges.

chmod +x createIRSA-AMPQuery.sh

3. Run the script.

Using Amazon Managed Service for Prometheus with interface VPC endpoints

If you use Amazon Virtual Private Cloud (Amazon VPC) to host your AWS resources, you can establish private connections between your VPC and Amazon Managed Service for Prometheus. You can use these connections to enable Amazon Managed Service for Prometheus to communicate with your resources on your VPC without going through the public internet. Amazon VPC is an AWS service that you can use to launch AWS resources in a virtual network that you define. With a VPC, you have control over your network settings, such the IP address range, subnets, route tables, and network gateways. To connect your VPC to Amazon Managed Service for Prometheus, you define an *interface VPC endpoint* to connect your VPC to AWS services. The endpoint provides reliable, scalable connectivity to Amazon Managed Service for Prometheus without requiring an internet gateway, a network address translation (NAT) instance, or a VPN connection. For more information, see <u>What Is Amazon VPC</u> in the *Amazon VPC User Guide*.

Interface VPC endpoints are powered by AWS PrivateLink, an AWS technology that enables private communication between AWS services using an elastic network interface with private IP addresses. For more information, see the <u>New – AWS PrivateLink for AWS Services</u> blog post.

The following information is for Amazon VPC users. For information about how to get started with Amazon VPC, see <u>Getting Started</u> in the *Amazon VPC User Guide*.

Create an interface VPC endpoint for Amazon Managed Service for Prometheus

Create an interface VPC endpoint to begin using Amazon Managed Service for Prometheus. Choose from the following service name endpoints:

com.amazonaws.region.aps-workspaces

Choose this service name to work with Prometheus-compatible APIs. For more information, see <u>Prometheus-compatible APIs</u> in the *Amazon Managed Service for Prometheus User Guide*.

com.amazonaws.region.aps

Choose this service name to perform workspace management tasks. For more information, see <u>Amazon Managed Service for Prometheus APIs</u> in the *Amazon Managed Service for Prometheus User Guide*.

🚯 Note

If you are using remote_write in a VPC without direct internet access, you must also create an interface VPC endpoint for AWS Security Token Service, to allow sigv4 to work through the endpoint. For information about creating a VPC endpoint for AWS STS, see <u>Using AWS</u> <u>STS interface VPC endpoints</u> in the AWS Identity and Access Management User Guide. You must set AWS STS to use regionalized endpoints.

For more information, including step-by-step instructions to create an interface VPC endpoint, see Creating an interface endpoint in the Amazon VPC User Guide.

Note

You can use **VPC endpoint policies** to control access to your Amazon Managed Service for Prometheus interface VPC endpoint. See the next section for more information.

If you created an interface VPC endpoint for Amazon Managed Service for Prometheus and already have data flowing to the workspaces located on your VPC, the metrics will flow through the interface VPC endpoint by default. Amazon Managed Service for Prometheus uses public endpoints or private interface endpoints (whichever are in use) to perform this task.

Controlling access to your Amazon Managed Service for Prometheus VPC endpoint

You can use VPC endpoint policies to control access to your Amazon Managed Service for Prometheus interface VPC endpoint. A VPC endpoint policy is an IAM resource policy that you attach to an endpoint when you create or modify the endpoint. If you don't attach a policy when you create an endpoint, Amazon VPC attaches a default policy for you that allows full access to the service. An endpoint policy doesn't override or replace IAM identity-based policies or servicespecific policies. It's a separate policy for controlling access from the endpoint to the specified service.

For more information, see <u>Controlling Access to Services with VPC Endpoints</u> in the *Amazon VPC User Guide*.

The following is an example of an endpoint policy for Amazon Managed Service for Prometheus. This policy allows users with the role PromUser connecting to Amazon Managed Service for Prometheus through the VPC to view workspaces and rule groups, but not, for example, to create or delete workspaces.

Create an interface VPC endpoint for Amazon Managed Service for Prometheus

```
"Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AmazonManagedPrometheusPermissions",
            "Effect": "Allow",
            "Action": [
                "aps:DescribeWorkspace",
                "aps:DescribeRuleGroupsNamespace",
                "aps:ListRuleGroupsNamespace",
                "aps:ListWorkspaces"
            ],
            "Resource": "arn:aws:aps:*:*:/workspaces*",
            "Principal": {
                "AWS": [
                     "arn:aws:iam::111122223333:role/PromUser"
                ]
            }
        }
    ]
}
```

The following example shows a policy that only allows requests coming from a specified IP address in the specified VPC to succeed. Requests from other IP addresses will fail.

```
{
   "Statement": [
       {
          "Action": "aps:*",
          "Effect": "Allow",
          "Principal": "*",
          "Resource": "*",
          "Condition": {
              "IpAddress": {
                 "aws:VpcSourceIp": "192.0.2.123"
              },
       "StringEquals": {
                 }
          }
       }
   ]
}
```

Troubleshoot Amazon Managed Service for Prometheus errors

Use the following sections to help troubleshoot issues with Amazon Managed Service for Prometheus.

Topics

- 429 or limit exceeded errors
- I see duplicate samples
- I see errors about sample timestamps
- I see an error message related to a limit
- Your local Prometheus server output exceeds the limit.
- Some of my data isn't appearing

429 or limit exceeded errors

If you see a 429 error similar to the following example, your requests have exceeded Amazon Managed Service for Prometheus ingestion quotas.

```
ts=2020-10-29T15:34:41.845Z caller=dedupe.go:112 component=remote level=error
remote_name=e13b0c
url=http://iamproxy-external.prometheus.uswest2-prod.eks:9090/workspaces/workspace_id/
api/v1/remote_write
msg="non-recoverable error" count=500 err="server returned HTTP status 429
Too Many Requests: ingestion rate limit (6666.6666666666666667) exceeded while adding 499
samples and 0 metadata
```

If you see a 429 error similar to the following example, your requests have exceeded the Amazon Managed Service for Prometheus quota for the number of active metrics in a workspace.

```
ts=2020-11-05T12:40:33.375Z caller=dedupe.go:112 component=remote level=error
remote_name=aps
url=http://iamproxy-external.prometheus.uswest2-prod.eks:9090/workspaces/workspace_id/
api/v1/remote_write
msg="non-recoverable error" count=500 err="server returned HTTP status 429 Too Many
Requests: user=accountid_workspace_id:
```

```
per-user series limit (local limit: 0 global limit: 3000000 actual local limit: 500000)
exceeded
```

If you see a 429 error similar to the following example, your requests have exceeded the Amazon Managed Service for Prometheus quota for the rate (transactions per second) that you can send data to your workspace using the RemoteWrite Prometheus compatible API.

```
ts=2024-03-26T16:50:21.780708811Z caller=dedupe.go:112 component=remote level=error
remote_name=ab123c
url=https://aps-workspaces.us-east-1.amazonaws.com/workspaces/workspace_id/api/v1/
remote_write
msg="non-recoverable error" count=1000 exemplarCount=0 err="server returned HTTP status
429 Too Many Requests: {\"message\":\"Rate exceeded\"}"
```

If you see a 400 error similar to the following example, your requests have exceeded Amazon Managed Service for Prometheus quota for active time series. For details about how active time series quotas are handled, see <u>Active series default</u>.

```
ts=2024-03-26T16:50:21.780708811Z caller=push.go:53 level=warn
url=https://aps-workspaces.us-east-1.amazonaws.com/workspaces/workspace_id/api/v1/
remote_write
msg="non-recoverable error" count=500 exemplarCount=0
err="server returned HTTP status 400 Bad Request: maxFailure (quorum) on a given error
family, rpc error: code = Code(400)
desc = addr=10.1.41.23:9095 state=ACTIVE zone=us-east-1a, rpc error: code = Code(400)
desc = user=accountid_workspace_id: per-user series limit of 10000000 exceeded,
Capacity from 2,000,000 to 10,000,000 is automatically adjusted based on the last 30
min of usage.
If throttled above 10,000,000 or in case of incoming surges, please contact
administrator to raise it.
(local limit: 0 global limit: 10000000 actual local limit: 92879)"
```

For more information about Amazon Managed Service for Prometheus service quotas and about how to request increases, see Amazon Managed Service for Prometheus service quotas

I see duplicate samples

If you are using a high-availability Prometheus group, you need to use external labels on your Prometheus instances to set up deduplication. For more information, see <u>Deduplicating high</u> availability metrics sent to Amazon Managed Service for Prometheus.

Other issues around duplicated data are discussed in the next section.

I see errors about sample timestamps

Amazon Managed Service for Prometheus ingests data in order, and expects each sample to have a timestamp later than the previous sample.

If your data does not arrive in order, you can see errors about out-of-order samples, duplicate sample for timestamp, or samples with different value but same timestamp. These issues are typically caused by incorrect setup of the client that is sending data to Amazon Managed Service for Prometheus. If you are using a Prometheus client running in agent mode, check the configuration for rules with duplicate series name, or duplicated targets. If your metrics provide the timestamp directly, check that they are not out of order.

For more details about how this works, or ways to check your setup, see the blog post <u>Understanding Duplicate Samples and Out-of-order Timestamp Errors in Prometheus</u> from *Prom Labs*.

I see an error message related to a limit

🚯 Note

Amazon Managed Service for Prometheus provides <u>CloudWatch usage metrics</u> to monitor Prometheus resource usage. Using the CloudWatch usage metrics alarm feature, you can monitor Prometheus resources and usage to prevent limit errors.

If you see one of the following error messages, you can request an increase in one of the Amazon Managed Service for Prometheus quotas to solve the issue. For more information, see <u>Amazon</u> Managed Service for Prometheus service quotas.

- per-user series limit of <value> exceeded, please contact administrator to raise it
- per-metric series limit of <*value*> exceeded, please contact administrator to raise it
- ingestion rate limit (...) exceeded
- series has too many labels (...) series: '%s'
- the query time range exceeds the limit (query length: xxx, limit: yyy)
- the query hit the max number of chunks limit while fetching chunks from ingesters

· Limit exceeded. Maximum workspaces per account.

Your local Prometheus server output exceeds the limit.

Amazon Managed Service for Prometheus has service quotas for the amount of data that a workspace can receive from Prometheus servers. To find the amount of data that your Prometheus server is sending to Amazon Managed Service for Prometheus, you can run the following queries on your Prometheus server. If you find that your Prometheus output is exceeding a Amazon Managed Service for Prometheus output is exceeding a Amazon Managed Service for Prometheus an increase of the corresponding service quota. For more information, see <u>Amazon Managed Service for Prometheus service quotas</u>.

Type of data	Query to use
Current active series	prometheu s_tsdb_he ad_series
Current ingestion rate	<pre>rate(prom etheus_ts db_head_s amples_ap pended_to tal[5m])</pre>
Most-to-least list of active series per metric name	<pre>sort_desc (count by(name) ({name! =""}))</pre>
Number of labels per metric series	group by(mylabe lname) ({name! =""})

Queries against your local self-run Prometheus server to find the output limits.

Some of my data isn't appearing

Data that is sent to Amazon Managed Service for Prometheus can be discarded for various reasons. The following table shows reasons that data might be discarded rather than being ingested.

You can track the amount and reasons that data is discarded using Amazon CloudWatch. For more information, see <u>Use CloudWatch metrics to monitor Amazon Managed Service for Prometheus</u> resources.

Reason	Meaning
greater_than_max_sample_age	Discarding log lines which are older than the current time
new-value-for-timestamp	Duplicate samples are sent with a different timestamp than was previously recorded
per_metric_series_limit	User has hit the active series per metric limit
per_user_series_limit	User has hit the total number of active series limit
rate_limited	Ingestion rate limited
sample-out-of-order	Samples are sent out of order and cannot be processed
label_value_too_long	Label value is longer than allowed character limit
max_label_names_per_series	User has hit the label names per metric
missing_metric_name	Metric name is not provided
metric_name_invalid	Invalid metric name provided
label_invalid	Invalid label provided
duplicate_label_names	Duplicate label names provided

Tagging in Amazon Managed Service for Prometheus

A *tag* is a custom attribute label that you or AWS assigns to an AWS resource. Each AWS tag has two parts:

- A *tag key* (for example, CostCenter, Environment, Project, or Secret). Tag keys are case sensitive.
- An optional field known as a *tag value* (for example, 111122223333, Production, or a team name). Omitting the tag value is the same as using an empty string. Like tag keys, tag values are case sensitive.

Together these are known as key-value pairs. You can have as many as 50 tags assigned to each workspace.

Tags help you identify and organize your AWS resources. Many AWS services support tagging, so you can assign the same tag to resources from different services to indicate that the resources are related. For example, you can assign the same tag to an Amazon Managed Service for Prometheus workspace that you assign to an Amazon S3 bucket. For more information about tagging strategies, see <u>Tagging AWS Resources</u>.

In Amazon Managed Service for Prometheus, both workspaces and rule groups namespaces can be tagged. You can use the console, the AWS CLI, APIs, or SDKs to add, manage, and remove tags for these resources. In addition to identifying, organizing, and tracking your workspaces and rule groups namespaces with tags, you can use tags in IAM policies to help control who can view and interact with your Amazon Managed Service for Prometheus resources.

Tag restrictions

The following basic restrictions apply to tags:

- Each resource can have a maximum of 50 tags.
- For each resource, each tag key must be unique, and each tag key can have only one value.
- The maximum tag key length is 128 Unicode characters in UTF-8.
- The maximum tag value length is 256 Unicode characters in UTF-8.
- If your tagging schema is used across multiple AWS services and resources, remember that other services may have restrictions on allowed characters. Generally allowed characters are letters, numbers, spaces representable in UTF-8, and the following characters: .: + = @ _ / (hyphen).

- Tag keys and values are case sensitive. As a best practice, decide on a strategy for capitalizing tags and consistently implement that strategy across all resource types. For example, decide whether to use Costcenter, costcenter, or CostCenter and use the same convention for all tags. Avoid using similar tags with inconsistent case treatment.
- Don't use aws:, AWS:, or any upper or lowercase combination of such as a prefix for either keys or values. These are reserved only for AWS use. You can't edit or delete tag keys or values with this prefix. Tags with this prefix do not count against your tags-per-resource limit.

Topics

- Tag Amazon Managed Service for Prometheus workspaces
- Tagging rule groups namespaces

Tag Amazon Managed Service for Prometheus workspaces

Tags are custom labels that can be assigned to a resource. They include a unique key and an optional value (in a key-value pair). Tags help you identify and organize your AWS resources. In Amazon Managed Service for Prometheus, workspaces (and rule groups namespaces) can be tagged. You can use the console, the AWS CLI, APIs, or SDKs to add, manage, and remove tags for these resources. In addition to identifying, organizing, and tracking your workspaces with tags, you can use tags in IAM policies to help control who can view and interact with your Amazon Managed Service for Prometheus.

Use the procedures in this section to work with tags for Amazon Managed Service for Prometheus workspaces.

Topics

- Add a tag to a workspace
- View tags for a workspace
- Edit tags for a workspace
- <u>Remove a tag from a workspace</u>

Add a tag to a workspace

Adding tags to an Amazon Managed Service for Prometheus workspace can help you identify and organize your AWS resources and manage access to them. First, you add one or more tags (key-

value pairs) to a workspace. After you have tags, you can create IAM policies to manage access to the workspace based on these tags. You can use the the console or the AWS CLI to add tags to an Amazon Managed Service for Prometheus workspace.

<u> Important</u>

Adding tags to a workspace can impact access to that workspace. Before you add a tag to a workspace, make sure to review any IAM policies that might use tags to control access to resources.

For more information about adding tags to an Amazon Managed Service for Prometheus workspace when you create it, see <u>Create a Amazon Managed Service for Prometheus workspace</u>.

Topics

- Add a tag to a workspace (console)
- Add a tag to a workspace (AWS CLI)

Add a tag to a workspace (console)

You can use the console to add one or more tags to a Amazon Managed Service for Prometheus workspace.

- 1. Open the Amazon Managed Service for Prometheus console at https://console.aws.amazon.com/prometheus/.
- 2. In the navigation pane, choose the menu icon.
- 3. Choose All workspaces.
- 4. Choose the workspace ID of the workspace that you want to manage.
- 5. Choose the **Tags** tab.
- 6. If no tags have been added to the Amazon Managed Service for Prometheus workspace, choose **Create tag**. Otherwise, choose **Manage tags**.
- 7. In **Key**, enter a name for the tag. You can add an optional value for the tag in **Value**.
- 8. (Optional) To add another tag, choose Add tag again.
- 9. When you have finished adding tags, choose **Save changes**.

Add a tag to a workspace (AWS CLI)

Follow these steps to use the AWS CLI to add a tag to an Amazon Managed Service for Prometheus workspace. To add a tag to a workspace when you create it, see <u>Create a Amazon Managed Service</u> for Prometheus workspace.

In these steps, we assume that you have already installed a recent version of the AWS CLI or updated to the current version. For more information, see <u>Installing the AWS Command Line</u> <u>Interface</u>.

At the terminal or command line, run the **tag-resource** command, specifying the Amazon Resource Name (ARN) of the workspace where you want to add tags and the key and value of the tag you want to add. You can add more than one tag to an workspace. For example, to tag an Amazon Managed Service for Prometheus workspace named **My-Workspace** with two tags, a tag key named *Status* with the tag value of *Secret*, and a tag key named *Team* with the tag value of *My*-*Team*:

```
aws amp tag-resource --resource-arn arn:aws:aps:us-
west-2:123456789012:workspaces/IDstring
--tags Status=Secret,Team=My-Team
```

If successful, this command returns nothing.

View tags for a workspace

Tags can help you identify and organize your AWS resources and manage access to them. For more information about tagging strategies, see <u>Tagging AWS Resources</u>.

View tags for an Amazon Managed Service for Prometheus workspace (console)

You can use the console to view the tags associated with a Amazon Managed Service for Prometheus workspace.

- 1. Open the Amazon Managed Service for Prometheus console at https://console.aws.amazon.com/prometheus/.
- 2. In the navigation pane, choose the menu icon.
- 3. Choose **All workspaces**.
- 4. Choose the workspace ID of the workspace that you want to manage.
- 5. Choose the **Tags** tab.

View tags for an Amazon Managed Service for Prometheus workspace (AWS CLI)

Follow these steps to use the AWS CLI to view the AWS tags for an workspace. If no tags have been added, the returned list is empty.

At the terminal or command line, run the **list-tags-for-resource** command. For example, to view a list of tag keys and tag values for a workspace:

```
aws amp list-tags-for-resource --resource-arn arn:aws:aps:us-
west-2:123456789012:workspace/IDstring
```

If successful, this command returns information similar to the following:

```
{
    "tags": {
        "Status": "Secret",
        "Team": "My-Team"
    }
}
```

Edit tags for a workspace

You can change the value for a tag associated with a workspace. You can also change the name of the key, which is equivalent to removing the current tag and adding a different one with the new name and the same value as the other key.

<u> Important</u>

Editing tags for an Amazon Managed Service for Prometheus workspace can impact access to that workspace. Before you edit the name (key) or value of a tag for a workspace, make sure to review any IAM policies that might use the key or value for a tag to control access to resources such as repositories.

Edit a tag for an Amazon Managed Service for Prometheus workspace (console)

You can use the console to edit the tags associated with a Amazon Managed Service for Prometheus workspace.

- Open the Amazon Managed Service for Prometheus console at <u>https://</u> console.aws.amazon.com/prometheus/.
- 2. In the navigation pane, choose the menu icon.
- 3. Choose All workspaces.
- 4. Choose the workspace ID of the workspace that you want to manage.
- 5. Choose the **Tags** tab.
- If no tags have been added to the workspace, choose Create tag. Otherwise, choose Manage tags.
- 7. In Key, enter a name for the tag. You can add an optional value for the tag in Value.
- 8. (Optional) To add another tag, choose Add tag again.
- 9. When you have finished adding tags, choose **Save changes**.

Edit tags for an Amazon Managed Service for Prometheus workspace (AWS CLI)

Follow these steps to use the AWS CLI to update a tag for a workspace. You can change the value for an existing key, or add another key.

At the terminal or command line, run the **tag-resource** command, specifying the Amazon Resource Name (ARN) of the Amazon Managed Service for Prometheus workspace where you want to update a tag and specify the tag key and tag value:

```
aws amp tag-resource --resource-arn arn:aws:aps:us-
west-2:123456789012:workspace/IDstring --tags Team=New-Team
```

Remove a tag from a workspace

You can remove one or more tags associated with a workspace. Removing a tag does not delete the tag from other AWS resources that are associated with that tag.

<u> Important</u>

Removing tags for a Amazon Managed Service for Prometheus workspace can impact access to that workspace. Before you remove a tag from a workspace, make sure to review any IAM policies that might use the key or value for a tag to control access to resources such as repositories.

Remove a tag from an Amazon Managed Service for Prometheus workspace (console)

You can use the console to remove the association between a tag and a workspace.

- 1. Open the Amazon Managed Service for Prometheus console at https://console.aws.amazon.com/prometheus/.
- 2. In the navigation pane, choose the menu icon.
- 3. Choose **All workspaces**.
- 4. Choose the workspace ID of the workspace that you want to manage.
- 5. Choose the **Tags** tab.
- 6. Choose Manage tags.
- 7. Find the tag that you want to delete, and choose **Remove**.

Remove a tag from an Amazon Managed Service for Prometheus workspace (AWS CLI)

Follow these steps to use the AWS CLI to remove a tag from an workspace. Removing a tag does not delete it, but simply removes the association between the tag and the workspace.

i Note

If you delete an Amazon Managed Service for Prometheus workspace, all tag associations are removed from the deleted workspace. You do not have to remove tags before you delete a workspace.

At the terminal or command line, run the **untag-resource** command, specifying the Amazon Resource Name (ARN) of the workspace where you want to remove tags and the tag key of the tag you want to remove. For example, to remove a tag on a workspace named **My-Workspace** with the tag key *Status*:

```
aws amp untag-resource --resource-arn arn:aws:aps:us-
west-2:123456789012:workspace/IDstring --tag-keys Status
```

If successful, this command returns nothing. To verify the tags associated with the workspace, run the **list-tags-for-resource** command.

Tagging rule groups namespaces

Tags are custom labels that can be assigned to a resource. They include a unique key and an optional value (in a key-value pair). Tags help you identify and organize your AWS resources. In Amazon Managed Service for Prometheus, rule groups namespaces (and workspaces) can be tagged. You can use the console, the AWS CLI, APIs, or SDKs to add, manage, and remove tags for these resources. In addition to identifying, organizing, and tracking your rule groups namespaces with tags, you can use tags in IAM policies to help control who can view and interact with your Amazon Managed Service for Prometheus resources.

Use the procedures in this section to work with tags for Amazon Managed Service for Prometheus rule groups namespaces.

Topics

- Add a tag to a rule groups namespace
- View tags for a rule groups namespace
- Edit tags for a rule groups namespace
- <u>Remove a tag from a rule groups namespace</u>

Add a tag to a rule groups namespace

Adding tags to an Amazon Managed Service for Prometheus rule groups namespaces can help you identify and organize your AWS resources and manage access to them. First, you add one or more tags (key-value pairs) to a rule groups namespace. After you have tags, you can create IAM policies to manage access to the namespace based on these tags. You can use the the console or the AWS CLI to add tags to an Amazon Managed Service for Prometheus rule groups namespace.

🔥 Important

Adding tags to a rule groups namespace can impact access to that rule groups namespace. Before you add a tag, make sure to review any IAM policies that might use tags to control access to resources.

For more information about adding tags to a rule groups namespace when you create it, see <u>Create</u> a rules file.

Topics

- Add a tag to a rule groups namespace (console)
- Add a tag to a rule groups namespace (AWS CLI)

Add a tag to a rule groups namespace (console)

You can use the console to add one or more tags to a Amazon Managed Service for Prometheus rule groups namespace.

- 1. Open the Amazon Managed Service for Prometheus console at <u>https://</u> console.aws.amazon.com/prometheus/.
- 2. In the navigation pane, choose the menu icon.
- 3. Choose **All workspaces**.
- 4. Choose the workspace ID of the workspace that you want to manage.
- 5. Choose the **Rules management** tab.
- 6. Choose the button next to the namespace name and choose **Edit**.
- 7. Choose Create tags, Add new tag.
- 8. In **Key**, enter a name for the tag. You can add an optional value for the tag in **Value**.
- 9. (Optional) To add another tag, choose Add new tag again.
- 10. When you have finished adding tags, choose **Save changes**.

Add a tag to a rule groups namespace (AWS CLI)

Follow these steps to use the AWS CLI to add a tag to an Amazon Managed Service for Prometheus rule groups namespace. To add a tag to a rule groups namespace when you create it, see <u>Upload a</u> rules configuration file to Amazon Managed Service for Prometheus.

In these steps, we assume that you have already installed a recent version of the AWS CLI or updated to the current version. For more information, see <u>Installing the AWS Command Line</u> <u>Interface</u>.

At the terminal or command line, run the **tag-resource** command, specifying the Amazon Resource Name (ARN) of the rule groups namespace where you want to add tags and the key and value of the tag you want to add. You can add more than one tag to an rule groups namespace. For example, to tag an Amazon Managed Service for Prometheus namespace named **My-Workspace** with two tags, a tag key named *Status* with the tag value of *Secret*, and a tag key named *Team* with the tag value of *My-Team*:

```
aws amp tag-resource \
    --resource-arn arn:aws:aps:us-
west-2:123456789012:rulegroupsnamespace/IDstring/namespace_name \
    --tags Status=Secret,Team=My-Team
```

If successful, this command returns nothing.

View tags for a rule groups namespace

Tags can help you identify and organize your AWS resources and manage access to them. For more information about tagging strategies, see <u>Tagging AWS Resources</u>.

View tags for an Amazon Managed Service for Prometheus rule groups namespace (console)

You can use the console to view the tags associated with a Amazon Managed Service for Prometheus rule groups namespace.

- 1. Open the Amazon Managed Service for Prometheus console at https://console.aws.amazon.com/prometheus/.
- 2. In the navigation pane, choose the menu icon.
- 3. Choose All workspaces.
- 4. Choose the workspace ID of the workspace that you want to manage.
- 5. Choose the **Rules management** tab.
- 6. Choose the namespace name.

View tags for an Amazon Managed Service for Prometheus workspace (AWS CLI)

Follow these steps to use the AWS CLI to view the AWS tags for a rule groups namespace. If no tags have been added, the returned list is empty.

At the terminal or command line, run the **list-tags-for-resource** command. For example, to view a list of tag keys and tag values for a rule groups namespace:

```
aws amp list-tags-for-resource --resource-arn rn:aws:aps:us-
west-2:123456789012:rulegroupsnamespace/IDstring/namespace_name
```

If successful, this command returns information similar to the following:

```
{
    "tags": {
        "Status": "Secret",
        "Team": "My-Team"
    }
}
```

Edit tags for a rule groups namespace

You can change the value for a tag associated with a rule groups namespace. You can also change the name of the key, which is equivalent to removing the current tag and adding a different one with the new name and the same value as the other key.

<u> Important</u>

Editing tags for an rule groups namespace can impact access to it. Before you edit the name (key) or value of a tag for a resource, make sure to review any IAM policies that might use the key or value for a tag to control access to resources.

Edit a tag for an Amazon Managed Service for Prometheus rule groups namespace (console)

You can use the console to edit the tags associated with a Amazon Managed Service for Prometheus rule groups namespace.

- Open the Amazon Managed Service for Prometheus console at <u>https://</u> console.aws.amazon.com/prometheus/.
- 2. In the navigation pane, choose the menu icon.
- 3. Choose All workspaces.
- 4. Choose the workspace ID of the workspace that you want to manage.
- 5. Choose the **Rules management** tab.

- 6. Choose the name of the namespace.
- 7. Choose Manage tags, Add new tag.
- 8. To change the value of an existing tag, enter the new value for **Value**.
- 9. o add an additional tag, choose Add new tag.
- 10. When you have finished adding and editing tags, choose **Save changes**.

Edit tags for an Amazon Managed Service for Prometheus rule groups namespace (AWS CLI)

Follow these steps to use the AWS CLI to update a tag for a rule groups namespace. You can change the value for an existing key, or add another key.

At the terminal or command line, run the **tag-resource** command, specifying the Amazon Resource Name (ARN) of the resource where you want to update a tag and specify the tag key and tag value:

```
aws amp tag-resource --resource-arn rn:aws:aps:us-
west-2:123456789012:rulegroupsnamespace/IDstring/namespace_name --tags Team=New-Team
```

Remove a tag from a rule groups namespace

You can remove one or more tags associated with a rule groups namespace. Removing a tag does not delete the tag from other AWS resources that are associated with that tag.

🔥 Important

Removing tags for a resource can impact access to that resource. Before you remove a tag from a resource, make sure to review any IAM policies that might use the key or value for a tag to control access to resources such as repositories.

Remove a tag from an Amazon Managed Service for Prometheus rule groups namespace (console)

You can use the console to remove the association between a tag and a rule groups namespace.

 Open the Amazon Managed Service for Prometheus console at <u>https://</u> console.aws.amazon.com/prometheus/.

- 2. In the navigation pane, choose the menu icon.
- 3. Choose All workspaces.
- 4. Choose the workspace ID of the workspace that you want to manage.
- 5. Choose the **Rules management** tab.
- 6. Choose the name of the namespace.
- 7. Choose Manage tags.
- 8. Next to the tag you want to delete, choose **Remove**.
- 9. When you have finished, choose **Save changes**.

Remove a tag from an Amazon Managed Service for Prometheus rule groups namespace (AWS CLI)

Follow these steps to use the AWS CLI to remove a tag from an rule groups namespace. Removing a tag does not delete it, but simply removes the association between the tag and the rule groups namespace.

🚺 Note

If you delete an Amazon Managed Service for Prometheus rule groups namespace, all tag associations are removed from the deleted nnamespace. You do not have to remove tags before you delete a namespace.

At the terminal or command line, run the **untag-resource** command, specifying the Amazon Resource Name (ARN) of the rule groups namespace where you want to remove tags and the tag key of the tag you want to remove. For example, to remove a tag on a workspace named **My-Workspace** with the tag key *Status*:

```
aws amp untag-resource --resource-arn rn:aws:aps:us-
west-2:123456789012:rulegroupsnamespace/IDstring/namespace_name --tag-keys Status
```

If successful, this command returns nothing. To verify the tags associated with the resource, run the **list-tags-for-resource** command.

Amazon Managed Service for Prometheus service quotas

The following two sections describe the quotas and limits associated with Amazon Managed Service for Prometheus.

Service quotas

Amazon Managed Service for Prometheus has the following quotas. Amazon Managed Service for Prometheus vends <u>CloudWatch usage metrics</u> to monitor Prometheus resource usage. Using the CloudWatch usage metrics alarm feature, you can monitor Prometheus resources and usage to prevent limit errors.

As your projects and workspaces grow, the most common quotas that you may need to monitor or request an increase for are: **Active series per workspace**, **Ingestion rate per workspace**, and **Ingestion burst size per workspace**.

For all adjustable quotas, you can request a quota increase by selecting the link in the **Adjustable** column, or by <u>requesting a quota increase</u>.

The Active series per workspace limit is dynamically applied. For more information, see <u>Active</u> <u>series default</u>. The *Ingestion rate per workspace* and *Ingestion burst size per workspace* together control how quickly you can ingest data into your workspace. For more information see <u>Ingestion throttling</u>.

🚺 Note

Unless otherwise noted, these quotas are per workspace. The maximum value for active series per workspace is one billion.

Name	Default	Adjus e	Description
Active metrics with metadata per workspace	Each supported Region: 20,000	No	The number of unique active metrics with metadata per workspace . Note: If the limit is

Name	Default	Adjus e	Description
			reached, metric sample is recorded, but metadata over the limit is dropped.
Active series per workspace	Each supported Region: 10,000,00 0 per 2 hours	Yes	The number of unique active series per workspace (up to a maximum of 1 billion). A series is active if a sample has been reported in the past 2 hours. Capacity from 2M to 10M is automatically adjusted based on the last 30 min of usage.
Alert aggregation group size in alert manager definition file	Each supported Region: 1,000	<u>Yes</u>	The maximum size of an alert aggregation group in alert manager definition file. Each label value combination of group_by would create an aggregation group.
Alert manager definition file size	Each supported Region: 1 Megabytes	No	The maximum size of an alert manager definition file.
Alert payload size in Alert Manager	Each supported Region: 20 Megabytes	No	The maximum alert payload size of all Alert Manager alerts per workspace. Alert size is dependent on labels and annotations.

Name	Default	Adjus e	Description
Alerts in Alert Manager	Each supported Region: 1,000	<u>Yes</u>	The maximum number of concurrent Alert Manager alerts per workspace.
HA tracker clusters	Each supported Region: 500	No	The maximum number of clusters that HA tracker will keep track of for ingested samples per workspace.
Ingestion burst size per workspace	Each supported Region: 1,000,000	<u>Yes</u>	The maximum number samples that could be ingested per workspace in one burst per second.
Ingestion rate per workspace	Each supported Region: 170,000	<u>Yes</u>	Metric sample ingestion rate per workspace per second.
Inhibition rules in alert manager definition file	Each supported Region: 100	<u>Yes</u>	The maximum number of inhibition rules in alert manager definition file.
Label size	Each supported Region: 7 Kilobytes	No	The maximum combined size of all labels and label values accepted for a series.
LabelSet limits per workspace	Each supported Region: 100	<u>Yes</u>	The maximum number of labelset limits that can be created per workspace

Name	Default	Adjus e	Description
Labels per metric series	Each supported Region: 150	<u>Yes</u>	Number of labels per metric series.
Metadata length	Each supported Region: 1 Kilobytes	No	The maximum length accepted for metric metadata. Metadata refers to Metric Name, Type, Unit and Help Text.
Metadata per metric	Each supported Region: 10	No	The number of metadata per metric. Note: If the limit is reached, metric sample is recorded, but metadata over the limit is dropped.
Nodes in alert manager routing tree	Each supported Region: 100	<u>Yes</u>	The maximum number of nodes in the alert manager routing tree.
Number of API operations per region in transactions per second	Each supported Region: 10	Yes	The maximum number of API operations per second per region for all Amazon Managed Service for Prometheu s APIs, including workspace CRUD APIs, tagging APIs, rule groups namespace CRUD APIs, and alert manager definition CRUD APIs.

Name	Default	Adjus e	Description
Number of GetSeries, GetLabels and GetMetricMetadata API operations per workspace in transactions per second	Each supported Region: 10	No	The maximum number of GetSeries, GetLabels and GetMetricMetadata Prometheus-compati ble API operations per second per workspace.
Number of QueryMetrics API operation s per workspace in transactions per second	Each supported Region: 300	No	The maximum number of QueryMetrics Prometheu s-compatible API operations per second per workspace.
Number of RemoteWrite API operation s per workspace in transactions per second	Each supported Region: 3,000	No	The maximum number of RemoteWrite Prometheu s-compatible API operations per second per workspace.
Number of other Prometheus-compati ble API operations per workspace in transactions per second	Each supported Region: 100	No	The maximum number of API operations per second per workspace for all other Prometheu s-compatible APIs including ListAlerts, ListRules, etc.
Query bytes for instant queries	Each supported Region: 5 Gigabytes	No	The maximum bytes that can be scanned by a single instant query.
Query bytes for range queries	Each supported Region: 5 Gigabytes	No	The maximum bytes that can be scanned per 24- hour interval in a single range query.

Name	Default	Adjus e	Description
Query chunks fetched	Each supported Region: 20,000,00 0	No	The maximum number of chunks that can be scanned during a single query.
Query samples	Each supported Region: 50,000,00 0	No	The maximum number of samples that can be scanned during a single query.
Query series fetched	Each supported Region: 12,000,00 0	No	The maximum number of series that can be scanned during a single query.
Query time range in days	Each supported Region: 95	No	The maximum time range of QueryMetrics, GetSeries, and GetLabels APIs.
Request size	Each supported Region: 1 Megabytes	No	The maximum request size for ingestion or query.
Rule evaluation interval	Each supported Region: 30 Seconds	<u>Yes</u>	The minimum rule evaluation interval of a rule group per workspace
Rule group namespace definition file size	Each supported Region: 1 Megabytes	No	The maximum size of a rule group namespace definition file.
Rules per workspace	Each supported Region: 2,000	<u>Yes</u>	The maximum number of rules per workspace.

Name	Default	Adjus e	Description
Silences per workspace	Each supported Region: 1,000	<u>Yes</u>	Maximum number of silences, including expired, active, and pending silences, per workspace.
Templates in alert manager definition file	Each supported Region: 100	<u>Yes</u>	The maximum number of templates in the alert manager definition file.
Workspaces per region per account	Each supported Region: 25	<u>Yes</u>	The maximum number of workspaces per region.

Active series default

Amazon Managed Service for Prometheus allows you to use up to your quota of active time series by default.

Amazon Managed Service for Prometheus workspaces automatically adapt to your ingestion volume. As your usage increases, Amazon Managed Service for Prometheus will automatically increase your time series capacity to double your baseline usage, up to the default quota. For example, if your average active time series for the last 30 minutes is 3.5 million, you can use up to 7 million time series without throttling.

If you need more than double your previous baseline, Amazon Managed Service for Prometheus automatically allocates more capacity as your ingest volume increases, to help ensure your workload does not experience sustained throttling, up to your quota. However, throttling can occur if you exceed double your previous baseline computed over the last 30 minutes. To avoid throttling, Amazon Managed Service for Prometheus recommends gradually increasing ingestion when increasing to more than double your previous active time series.

🚯 Note

The minimum capacity for active time series is 2 million, there is no throttling when you have less than 2 million series.

Ingestion throttling

Amazon Managed Service for Prometheus throttles ingestion for each workspace, based on your current limits. This helps maintain the performance of the workspace. If you exceed the limit, you will see DiscardedSamples in CloudWatch metrics (with the rate_limited reason). You can use Amazon CloudWatch to monitor your ingestion, and to create an alarm to warn you when you are close to reaching the throttling limits. For more information, see <u>Use CloudWatch metrics to monitor Amazon Managed Service for Prometheus resources</u>.

Amazon Managed Service for Prometheus uses the <u>token bucket algorithm</u> to implement ingestion throttling. With this algorithm, your account has a *bucket* that holds a specific number of *tokens*. The number of tokens in the bucket represents your ingestion limit at any given second.

Each data sample ingested removes one token from the bucket. If your bucket size (*Ingestion burst size per workspace*) is *1,000,000*, your workspace can ingest one million data samples in one second. If it exceeds one million samples to ingest, it will be throttled, and will not ingest any more records. Additional data samples will be discarded.

The bucket automatically refills at a set rate. If the bucket is below its maximum capacity, a set number of tokens is added back to it every second until it reaches its maximum capacity. If the bucket is full when the refill tokens arrive, they are discarded. The bucket can't hold more than its maximum number of tokens. The refill rate for sample ingestion is set by the *Ingestion rate per workspace* limit. If your *Ingestion rate per workspace* is set to 170,000, then the refill rate for the bucket is 170,000 tokens per second.

If your workspace ingests 1,000,000 data samples in a second, your bucket is immediately reduced to zero tokens. The bucket is then refilled by 170,000 tokens every second, until it reaches it's maximum capacity of 1,000,000 tokens. If there is no more ingestion, the previously empty bucket will return to it's maximum capacity in 6 seconds.

🚯 Note

Ingestion happens in batched requests. If you have 100 tokens available, and send a request with 101 samples, the entire request is rejected. Amazon Managed Service for

Prometheus does not partially accept requests. If you are writing a collector, you can manage retries (with smaller batches or after some time has passed).

You do not need to wait for the bucket to be full before your workspace can ingest more data samples. You can use tokens as they are added to the bucket. If you immediately use the refill tokens, the bucket does not reach its maximum capacity. For example, if you deplete the bucket, you can continue to ingest 170,000 data samples per second. The bucket can refill to maximum capacity only if you ingest fewer than 170,000 data samples per second.

Additional limits on ingested data

Amazon Managed Service for Prometheus also has the following additional requirements for data ingested into the workspace. These are not adjustable.

- Metric samples older than 1 hour are refused from being ingested.
- Every sample and metadata must have a metric name.

Amazon Managed Service for Prometheus API Reference

Amazon Managed Service for Prometheus offers two types of APIs:

- Amazon Managed Service for Prometheus APIs These APIs allow you to create and manage your Amazon Managed Service for Prometheus workspaces, including operations for workspaces, scrapers, alert manager definitions, rule groups namespaces, and logging. You use the AWS SDKs, available for various programming languages, to interact with these APIs.
- 2. **Prometheus-compatible APIs** Amazon Managed Service for Prometheus supports HTTP APIs that are compatible with Prometheus. These APIs enable building custom applications, automate workflows, integrate with other services or tools, and query and interact with your monitoring data using the Prometheus query language (PromQL).

This section lists the API operations and data structures supported by Amazon Managed Service for Prometheus.

For information about quotas for the series, labels, and API requests, see <u>Amazon Managed Service</u> for Prometheus service quotas in the Amazon Managed Service for Prometheus User Guide.

Topics

- Amazon Managed Service for Prometheus APIs
- Prometheus-compatible APIs

Amazon Managed Service for Prometheus APIs

Amazon Managed Service for Prometheus provides API operations creating and maintaining your Amazon Managed Service for Prometheus workspaces. This includes APIs for workspaces, scrapers, alert manager definitions, rule groups namespaces, and logging.

For detailed information about the Amazon Managed Service for Prometheus APIs, see the <u>Amazon</u> Managed Service for Prometheus API Reference.

Using Amazon Managed Service for Prometheus with an AWS SDK

AWS software development kits (SDKs) are available for many popular programming languages. Each SDK provides an API, code examples, and documentation that makes it easier for developers to build AWS applications in their preferred language. For a list of SDKs and tools by language, see Tools to Build on AWS in the AWS *Developer Center*.

SDK Versions

We recommend that you use the most recent build of the AWS SDK, and any other SDKs, that you use in your projects, and to keep the SDKs up to date. The AWS SDK provides you with the latest features and functionality, and also security updates.

Prometheus-compatible APIs

Amazon Managed Service for Prometheus supports the following Prometheus-compatible APIs.

For more information about using Prometheus-compatible APIs, see <u>Query using Prometheus-</u> compatible APIs.

Topics

- <u>CreateAlertManagerAlerts</u>
- DeleteAlertManagerSilence
- GetAlertManagerStatus
- GetAlertManagerSilence
- GetLabels
- GetMetricMetadata
- GetSeries
- ListAlerts
- ListAlertManagerAlerts
- ListAlertManagerAlertGroups
- ListAlertManagerReceivers
- ListAlertManagerSilences
- ListRules
- PutAlertManagerSilences
- QueryMetrics
- <u>RemoteWrite</u>

CreateAlertManagerAlerts

The CreateAlertManagerAlerts operation creates an alert in the workspace.

Valid HTTP verbs:

POST

Valid URIs:

/workspaces/workspaceId/alertmanager/api/v2/alerts

URL query parameters:

alerts An array of objects, where each object represents one alert. The following is an example of an alert object:

```
Ε
  {
    "startsAt": "2021-09-24T17:14:04.995Z",
    "endsAt": "2021-09-24T17:14:04.995Z",
    "annotations": {
      "additionalProp1": "string",
      "additionalProp2": "string",
      "additionalProp3": "string"
    },
    "labels": {
      "additionalProp1": "string",
      "additionalProp2": "string",
      "additionalProp3": "string"
    },
    "generatorURL": "string"
  }
]
```

Sample request

POST /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/alertmanager/api/v2/alerts
HTTP/1.1
Content-Length: 203,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0

```
[
    {
        "labels": {
            "alertname": "test-alert"
        },
        "annotations": {
            "summary": "this is a test alert used for demo purposes"
        },
        "generatorURL": "https://www.amazon.com/"
    }
]
```

Sample response

```
HTTP/1.1 200 OK
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Content-Length: 0
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
vary: Origin
```

DeleteAlertManagerSilence

The DeleteSilence deletes one alert silence.

Valid HTTP verbs:

DELETE

Valid URIs:

/workspaces/workspaceId/alertmanager/api/v2/silence/silenceID

URL query parameters: none

Sample request

```
DELETE /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/alertmanager/api/v2/silence/
d29d9df3-9125-4441-912c-70b05f86f973 HTTP/1.1
```

Content-Length: 0, Authorization: AUTHPARAMS X-Amz-Date: 20201201T193725Z User-Agent: Grafana/8.1.0

Sample response

```
HTTP/1.1 200 OK
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Content-Length: 0
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
vary: Origin
```

GetAlertManagerStatus

The GetAlertManagerStatus retrieves information about the status of alert manager.

Valid HTTP verbs:

GET

Valid URIs:

/workspaces/workspaceId/alertmanager/api/v2/status

URL query parameters: none

Sample request

```
GET /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/alertmanager/api/v2/status
HTTP/1.1
Content-Length: 0,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0
```

Sample response

HTTP/1.1 200 OK

```
User Guide
```

```
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Content-Length: 941
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
vary: Origin
{
    "cluster": null,
    "config": {
        "original": "global:\n resolve_timeout: 5m\n http_config:\n
 follow_redirects: true\n smtp_hello: localhost\n smtp_require_tls: true\nroute:
\n receiver: sns-0\n group_by:\n - label\n continue: false\nreceivers:\n-
 name: sns-0\n sns_configs:\n - send_resolved: false\n
                                                            http_config:\n
      follow_redirects: true\n
                                  sigv4: \{\}\n
                                                 topic_arn: arn:aws:sns:us-
                            subject: '{{ template \"sns.default.subject\" . }}'\n
west-2:123456789012:test\n
    message: '{{ template \"sns.default.message\" . }}'\n
                                                             workspace_arn:
 arn:aws:aps:us-west-2:123456789012:workspace/ws-58a6a446-5ec4-415b-9052-a449073bbd0a
\ntemplates: []\n"
    },
    "uptime": null,
    "versionInfo": null
}
```

GetAlertManagerSilence

The GetAlertManagerSilence retrieves information about one alert silence.

Valid HTTP verbs:

GET

Valid URIs:

/workspaces/workspaceId/alertmanager/api/v2/silence/silenceID

URL query parameters: none

Sample request

```
GET /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/alertmanager/api/v2/silence/
d29d9df3-9125-4441-912c-70b05f86f973 HTTP/1.1
```

```
Content-Length: 0,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0
```

Sample response

```
HTTP/1.1 200 OK
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Content-Length: 310
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
vary: Origin
{
    "id": "d29d9df3-9125-4441-912c-70b05f86f973",
    "status": {
        "state": "active"
    },
    "updatedAt": "2021-10-22T19:32:11.763Z",
    "comment": "hello-world",
    "createdBy": "test-person",
    "endsAt": "2023-07-24T01:05:36.000Z",
    "matchers": [
        {
            "isEqual": true,
            "isRegex": true,
            "name": "job",
            "value": "hello"
        }
    ],
    "startsAt": "2021-10-22T19:32:11.763Z"
}
```

GetLabels

The GetLabels operation retrieves the labels associated with a time series.

Valid HTTP verbs:

GET, POST

Valid URIs:

```
/workspaces/workspaceId/api/v1/labels
```

/workspaces/workspaceId/api/v1/label/label-name/values This URI supports only GET requests.

URL query parameters:

match[]=<series_selector> Repeated series selector argument that selects the series from which to read the label names. Optional.

```
start=<rfc3339 | unix_timestamp> Start timestamp. Optional.
```

end=<rfc3339 | unix_timestamp> End timestamp. Optional.

Sample request for /workspaces/workspaceId/api/v1/labels

```
GET /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/api/v1/labels HTTP/1.1
Content-Length: 0,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0
```

Sample response for /workspaces/workspaceId/api/v1/labels

```
HTTP/1.1 200 OK
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Content-Length: 1435
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
vary: Origin
{
    "status": "success",
    "data": [
        "___name___",
        "access_mode",
        "address",
        "alertname",
```

]

}

```
"alertstate",
"apiservice",
"app",
"app_kubernetes_io_instance",
"app_kubernetes_io_managed_by",
"app_kubernetes_io_name",
"area",
"beta_kubernetes_io_arch",
"beta_kubernetes_io_instance_type",
"beta_kubernetes_io_os",
"boot_id",
"branch",
"broadcast",
"buildDate",
....
```

Sample request for /workspaces/workspaceId/api/v1/label/label-name/values

GET /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/api/v1/label/access_mode/values
HTTP/1.1
Content-Length: 0,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0

Sample response for /workspaces/workspaceId/api/v1/label/label-name/values

```
HTTP/1.1 200 OK
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Content-Length: 74
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
vary: Origin
{
    "status": "success",
    "data": [
        "ReadWriteOnce"
]
```

}

GetMetricMetadata

The GetMetricMetadata operation retrieves metadata about metrics that are currently being scraped from targets. It does not provide any target information.

The data section of the query result consists of an object where each key is a metric name and each value is a list of unique metadata objects, as exposed for that metric name across all targets.

Valid HTTP verbs:

GET

Valid URIs:

```
/workspaces/workspaceId/api/v1/metadata
```

URL query parameters:

limit=<number> The maximum number of metrics to return.

metric=<string> A metric name to filter metadata for. If you keep this empty, all metric metadata is retrieved.

Sample request

```
GET /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/api/v1/metadata HTTP/1.1
Content-Length: 0,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0
```

```
HTTP/1.1 200 OK
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
Transfer-Encoding: chunked
```

GetSeries

The GetSeries operation retrieves list of time series that match a certain label set.

Valid HTTP verbs:

GET, POST

Valid URIs:

/workspaces/workspaceId/api/v1/series

URL query parameters:

```
match[]=<series_selector> Repeated series selector argument that selects the series to
return. At least one match[] argument must be provided.
```

start=<rfc3339 | unix_timestamp> Start timestamp. Optional

end=<rfc3339 | unix_timestamp> End timestamp. Optional

Sample request

```
POST /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/api/v1/series --data-urlencode
'match[]=node_cpu_seconds_total{app="prometheus"}' --data-urlencode 'start=1634936400'
--data-urlencode 'end=1634939100' HTTP/1.1
Content-Length: 0,
```

```
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0
```

```
HTTP/1.1 200 OK
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
content-encoding: gzip
{
    "status": "success",
    "data": [
        {
            "___name___": "node_cpu_seconds_total",
            "app": "prometheus",
            "app_kubernetes_io_managed_by": "Helm",
            "chart": "prometheus-11.12.1",
            "cluster": "cluster-1",
            "component": "node-exporter",
            "cpu": "0",
            "heritage": "Helm",
            "instance": "10.0.100.36:9100",
            "job": "kubernetes-service-endpoints",
            "kubernetes_name": "servicesstackprometheuscf14a6d7-node-exporter",
            "kubernetes_namespace": "default",
            "kubernetes_node": "ip-10-0-100-36.us-west-2.compute.internal",
            "mode": "idle",
            "release": "servicesstackprometheuscf14a6d7"
        },
        {
            "___name___": "node_cpu_seconds_total",
            "app": "prometheus",
            "app_kubernetes_io_managed_by": "Helm",
            "chart": "prometheus-11.12.1",
            "cluster": "cluster-1",
            "component": "node-exporter",
            "cpu": "0",
            "heritage": "Helm",
```

```
"instance": "10.0.100.36:9100",
"job": "kubernetes-service-endpoints",
"kubernetes_name": "servicesstackprometheuscf14a6d7-node-exporter",
"kubernetes_namespace": "default",
"kubernetes_node": "ip-10-0-100-36.us-west-2.compute.internal",
"mode": "iowait",
"release": "servicesstackprometheuscf14a6d7"
},
...
```

ListAlerts

}

The ListAlerts operation retrieves currently active alerts in the workspace.

Valid HTTP verbs:

GET

Valid URIs:

```
/workspaces/workspaceId/api/v1/alerts
```

Sample request

```
GET /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/api/v1/alerts HTTP/1.1
Content-Length: 0,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0
```

```
HTTP/1.1 200 OK
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Content-Length: 386
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
vary: Origin
```

```
{
  "status": "success",
  "data": {
    "alerts": [
      {
        "labels": {
          "alertname": "test-1.alert",
          "severity": "none"
        },
        "annotations": {
          "message": "message"
        },
        "state": "firing",
        "activeAt": "2020-12-01T19:37:25.429565909Z",
        "value": "1e+00"
      }
    ]
  },
  "errorType": "",
  "error": ""
}
```

ListAlertManagerAlerts

The ListAlertManagerAlerts retrieves information about the alerts currently firing in alert manager in the workspace.

Valid HTTP verbs:

GET

Valid URIs:

/workspaces/workspaceId/alertmanager/api/v2/alerts

Sample request

```
GET /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/alertmanager/api/v2/alerts
HTTP/1.1
Content-Length: 0,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
```

User-Agent: Grafana/8.1.0

Sample response

```
HTTP/1.1 200 OK
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Content-Length: 354
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
vary: Origin
Ε
    {
        "annotations": {
            "summary": "this is a test alert used for demo purposes"
        },
        "endsAt": "2021-10-21T22:07:31.501Z",
        "fingerprint": "375eab7b59892505",
        "receivers": [
            {
                "name": "sns-0"
            }
        ],
        "startsAt": "2021-10-21T22:02:31.501Z",
        "status": {
            "inhibitedBy": [],
            "silencedBy": [],
            "state": "active"
        },
        "updatedAt": "2021-10-21T22:02:31.501Z",
        "labels": {
            "alertname": "test-alert"
        }
    }
]
```

ListAlertManagerAlertGroups

The ListAlertManagerAlertGroups operation retrieves a list of alert groups configured in alert manager in the workspace.

Valid HTTP verbs:

GET

Valid URIs:

/workspaces/workspaceId/alertmanager/api/v2/alerts/groups

URL query parameters:

active Boolean. If true, the returned list includes active alerts. The default is true. Optional

silenced Boolean. If true, the returned list includes silenced alerts. The default is true. Optional

inhibited Boolean. If true, the returned list includes inhibited alerts. The default is true. Optional

filter An array of strings. A list of matchers to filter alerts by. Optional

receiver String. A regular expression matching receivers to filter alerts by. Optional

Sample request

```
GET /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/alertmanager/api/v2/alerts/
groups HTTP/1.1
Content-Length: 0,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0
```

```
HTTP/1.1 200 OK
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Content-Length: 443
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
vary: Origin
[
{
```

```
"alerts": [
            {
                 "annotations": {
                     "summary": "this is a test alert used for demo purposes"
                },
                "endsAt": "2021-10-21T22:07:31.501Z",
                "fingerprint": "375eab7b59892505",
                "receivers": [
                     {
                         "name": "sns-0"
                     }
                ],
                "startsAt": "2021-10-21T22:02:31.501Z",
                 "status": {
                     "inhibitedBy": [],
                     "silencedBy": [],
                     "state": "unprocessed"
                },
                "updatedAt": "2021-10-21T22:02:31.501Z",
                "generatorURL": "https://www.amazon.com/",
                 "labels": {
                     "alertname": "test-alert"
                }
            }
        ],
        "labels": {},
        "receiver": {
            "name": "sns-0"
        }
    }
]
```

ListAlertManagerReceivers

The ListAlertManagerReceivers operation retrieves information about the receivers configured in alert manager.

Valid HTTP verbs:

GET

Valid URIs:

/workspaces/workspaceId/alertmanager/api/v2/receivers

URL query parameters: none

Sample request

```
GET /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/alertmanager/api/v2/receivers
HTTP/1.1
Content-Length: 0,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0
```

Sample response

ListAlertManagerSilences

The ListAlertManagerSilences operation retrieves information about the alert silences configured in the workspace.

Valid HTTP verbs:

GET

Valid URIs:

/workspaces/workspaceId/alertmanager/api/v2/silences

Sample request

```
GET /workspaces/ws-58a6a446-5ec4-415b-9052-a449073bbd0a/alertmanager/api/v2/silences
HTTP/1.1
Content-Length: 0,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0
```

```
HTTP/1.1 200 OK
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Content-Length: 312
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
vary: Origin
Γ
    {
        "id": "d29d9df3-9125-4441-912c-70b05f86f973",
        "status": {
            "state": "active"
        },
        "updatedAt": "2021-10-22T19:32:11.763Z",
        "comment": "hello-world",
        "createdBy": "test-person",
        "endsAt": "2023-07-24T01:05:36.000Z",
        "matchers": [
            {
                "isEqual": true,
                "isRegex": true,
                "name": "job",
                "value": "hello"
            }
        ],
        "startsAt": "2021-10-22T19:32:11.763Z"
    }
]
```

ListRules

The ListRules retrieves information about the rules configured in the workspace.

Valid HTTP verbs:

GET

Valid URIs:

/workspaces/workspaceId/api/v1/rules

Sample request

```
GET /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/api/v1/rules HTTP/1.1
Content-Length: 0,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0
```

```
HTTP/1.1 200 OK
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Content-Length: 423
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
vary: Origin
{
    "status": "success",
    "data": {
        "groups": [
            {
                "name": "test-1.rules",
                "file": "test-rules",
                "rules": [
                    {
                        "name": "record:1",
                        "query": "sum(rate(node_cpu_seconds_total[10m:1m]))",
```

```
"labels": {},
                         "health": "ok",
                         "lastError": "",
                         "type": "recording",
                         "lastEvaluation": "2021-10-21T21:22:34.429565909Z",
                         "evaluationTime": 0.001005399
                     }
                ],
                "interval": 60,
                "lastEvaluation": "2021-10-21T21:22:34.429563992Z",
                "evaluationTime": 0.001010504
            }
        ]
    },
    "errorType": "",
    "error": ""
}
```

PutAlertManagerSilences

The PutAlertManagerSilences operation creates a new alert silence or updates an existing one.

Valid HTTP verbs:

POST

Valid URIs:

```
/workspaces/workspaceId/alertmanager/api/v2/silences
```

URL query parameters:

silence An object that represents the silence. The following is the format:

```
{
  "id": "string",
  "matchers": [
    {
        "name": "string",
        "value": "string",
        "isRegex": Boolean,
        "isEqual": Boolean
}
```

```
],
"startsAt": "timestamp",
"endsAt": "timestamp",
"createdBy": "string",
"comment": "string"
}
```

Sample request

```
POST /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/alertmanager/api/v2/silences
 HTTP/1.1
Content-Length: 281,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0
{
   "matchers":[
      {
         "name":"job",
         "value":"up",
         "isRegex":false,
         "isEqual":true
      }
   ],
   "startsAt":"2020-07-23T01:05:36+00:00",
   "endsAt":"2023-07-24T01:05:36+00:00",
   "createdBy":"test-person",
   "comment":"test silence"
}
```

Sample response

HTTP/1.1 200 OK x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535 Content-Length: 53 Connection: keep-alive Date: Tue, 01 Dec 2020 19:37:25 GMT Content-Type: application/json Server: amazon vary: Origin "silenceID": "512860da-74f3-43c9-8833-cec026542b32"

QueryMetrics

{

}

The QueryMetrics operation evaluates an instant query at a single point in time or over a range of time.

Valid HTTP verbs:

GET, POST

Valid URIs:

/workspaces/workspaceId/api/v1/query This URI evaluates an instant query at a single point in time.

/workspaces/workspaceId/api/v1/query_range This URI evaluates an instant query over a range of time.

URL query parameters:

query=<string> A Prometheus expression query string. Used in both query and query_range.

time=<rfc3339 | unix_timestamp> (Optional) Evaluation timestamp if you are using the query for an instant query at a single point in time.

timeout=<duration> (Optional) Evaluation timeout. Defaults to and is capped by the value
of the -query.timeout flag. Used in both query and query_range.

start=<rfc3339 | unix_timestamp> Start timestamp if you are using query_range to query for a range of time.

end=<rfc3339 | unix_timestamp> End timestamp if you are using query_range to query for a range of time.

step=<duration | float> Query resolution step width in duration format or as a float
number of seconds. Use only if you are using query_range to query for a range of time, and
required for such queries.

max_samples_processed_warning_threshold=<integer> (Optional) Sets the warning threshold for Query Samples Processed (QSP). When queries hit this threshold, a warning message will be returned in the API response.

max_samples_processed_error_threshold=<integer>> (Optional) Sets the error threshold for Query Samples Processed (QSP). Queries that exceed this threshold will be rejected with an error and will not be charged. Used to prevent excessive query costs.

Duration

A duration in a Prometheus-compatible API is a number, followed immediately by one of the following units:

- ms milliseconds
- s seconds
- m minutes
- h hours
- d days, assuming a day always has 24h
- w weeks, assuming a week always has 7d
- y years, assuming a year always has 365d

Sample request

```
POST /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/api/v1/query?
query=sum(node_cpu_seconds_total) HTTP/1.1
Content-Length: 0,
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Grafana/8.1.0
```

Sample response

HTTP/1.1 200 OK x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535 Content-Length: 132 Connection: keep-alive Date: Tue, 01 Dec 2020 19:37:25 GMT

```
Content-Type: application/json
Server: amazon
content-encoding: gzip
{
    "status": "success",
    "data": {
        "resultType": "vector",
        "result": [
            {
                 "metric": {},
                 "value": [
                     1634937046.322,
                     "252590622.81000024"
                 ]
            }
        ]
    }
}
```

RemoteWrite

The RemoteWrite operation writes metrics from a Prometheus server to a remote URL in a standardized format. Typically, you will use an existing client such as a Prometheus server to call this operation.

Valid HTTP verbs:

POST

Valid URIs:

/workspaces/workspaceId/api/v1/remote_write

URL query parameters:

None

RemoteWrite has an ingestion rate of 70,000 samples per second and ingestion burst size of 1,000,000 samples.

Sample request

```
POST /workspaces/ws-b226cc2a-a446-46a9-933a-ac50479a5568/api/v1/remote_write --data-
binary "@real-dataset.sz" HTTP/1.1
Authorization: AUTHPARAMS
X-Amz-Date: 20201201T193725Z
User-Agent: Prometheus/2.20.1
Content-Type: application/x-protobuf
Content-Encoding: snappy
X-Prometheus-Remote-Write-Version: 0.1.0
```

body

Note

For the request body syntax, see to the protocol buffer definition at <u>https://github.com/</u> prometheus/prometheus/blob/1c624c58ca934f618be737b4995e22051f5724c1/prompb/ remote.pb.go#L64.

```
HTTP/1.1 200 OK
x-amzn-RequestId: 12345678-abcd-4442-b8c5-262b45e9b535
Content-Length:0
Connection: keep-alive
Date: Tue, 01 Dec 2020 19:37:25 GMT
Content-Type: application/json
Server: amazon
vary: Origin
```

Document History for Amazon Managed Service for Prometheus User Guide

The following table describes important documentation updates in the Amazon Managed Service for Prometheus User Guide. For notification about updates to this documentation, you can subscribe to an RSS feed.

Change	Description	Date
Update to the AmazonPro metheusConsoleFullAccess managed IAM policy.	The AmazonPrometheusCo nsoleFullAccess policy was updated. The aps:Creat eQueryLoggingConfi guration , aps:Updat eQueryLoggingConfi guration , aps:Delet eQueryLoggingConfi guration , aps:Descr ibeQueryLoggingCon figuration permissions were added to the policy.	May 5, 2025
Added editing of rules definition files and Alert manager configuration files in the console	Amazon Managed Service for Prometheus adds support for editing <u>Alert manager</u> <u>configuration files</u> and <u>rules</u> <u>definition files</u> from within the Amazon Managed Service for Prometheus console.	May 16, 2024
Added simpler AWS managed collector setup with access entries for Amazon EKS	Amazon Managed Service for Prometheus adds support for Amazon EKS access entries to simplify setting up <u>AWS managed collector</u> <u>s</u> . The <u>AmazonPrometheusSc</u>	May 2, 2024

	raperServiceRolePolicy managed policy for AWS managed collectors is updated to allow deleting access entries that are no longer used.	
<u>Move AWS API to a separate</u> <u>API reference guide</u>	The Amazon Managed Service for Prometheus AWS APIs are now available in their own reference, the <u>Amazon Managed Service for</u> <u>Prometheus API Reference</u> . Prometheus-compatible APIs continue to be documented in the <u>Amazon Managed Service</u> <u>for Prometheus User Guide</u> .	February 7, 2024
<u>Added customer managed</u> <u>keys for workspace encryptio</u> <u>n</u>	Amazon Managed Service for Prometheus adds support for customer managed keys for workspace encryption. For more information, see <u>Encryption at rest</u> .	December 21, 2023
Added new permissions to AmazonPrometheusFu llAccess	Added new permissions to the <u>AmazonPrometheusFu</u> <u>llAccess</u> managed policy to support creating AWS managed collectors for	November 26, 2023

Amazon EKS clusters.

Added new managed policy, AmazonPrometheusSc raperServiceLinkedRolePolicy	Added a new managed policy, <u>AmazonPrometheusSc</u> <u>raperServiceLinkedRolePolic</u> y for AWS managed collector s to collect metrics from Amazon EKS clusters.	November 26, 2023
Added AWS managed collectors as ingestion method	Amazon Managed Service for Prometheus adds support for <u>AWS managed collectors</u> .	November 26, 2023
Added support for integrati ng with Amazon Managed Grafana	Amazon Managed Service for Prometheus adds support for <u>integrating with Amazon</u> <u>Managed Grafana alerts</u> .	November 23, 2022
Added new permissions to AmazonPrometheusCo nsoleFullAccess	Added new permissions to the <u>AmazonPrometheusCo</u> <u>nsoleFullAccess</u> managed policy to support logging alert manager and ruler events in CloudWatch Logs.	October 24, 2022
Added Amazon EKS observabi lity solution.	Amazon Managed Service for Prometheus adds a new solution using AWS Observabi lity Accelerator. For more information, see <u>Using AWS</u> <u>Observability Accelerator</u> .	October 14, 2022
Added support for integrati ng into Amazon EKS cost monitoring.	Amazon Managed Service for Prometheus adds support for integrating into Amazon EKS cost monitoring. For more information, see <u>Integrati</u> <u>ng with Amazon EKS cost</u> <u>monitoring</u> .	September 22, 2022

Launched support for Alert Manager and Ruler logs in Amazon CloudWatch Logs.	Amazon Managed Service for Prometheus launches support for Alert Manager and Ruler error logs in Amazon CloudWatch Logs. For more information, see <u>Amazon</u> <u>CloudWatch Logs</u> .	September 1, 2022
<u>Added custom storage</u> <u>retention support.</u>	Amazon Managed Service for Prometheus adds custom storage retention support, per workspace, by modifying the quota for that workspace . For more information about quotas in Amazon Managed Service for Prometheus, see <u>Service quotas</u> .	August 12, 2022
Added usage metrics to Amazon CloudWatch.	Amazon Managed Service for Prometheus adds support for sending usage metrics to Amazon CloudWatch. For more information, see <u>Amazon CloudWatch metrics</u> .	May 6, 2022
<u>Added support for the Europe</u> (London) Region.	Amazon Managed Service for Prometheus adds support for the Europe (London) Region.	May 4, 2022
Amazon Managed Service for Prometheus is generally available, and adds support for rules and alert manager.	Amazon Managed Service for Prometheus is generally available. It also supports rules and alert manager. For more information, see <u>Recording rules and alerting</u> <u>rules</u> and <u>Alert manager and</u> <u>templating</u> .	September 29, 2021

Tagging support added.	Amazon Managed Service for Prometheus supports tagging of Amazon Managed Service for Prometheus workspaces.	September 7, 2021
Active series and ingestion rate quotas increased.	The active series quota increased to 1,000,000 and the ingestion rate quota increased to 70,000 samples per second.	February 22, 2021
Amazon Managed Service for Prometheus preview release.	The preview of Amazon Managed Service for Prometheus is released.	December 15, 2020