Amazon Managed Service for Prometheus pricing

Why Amazon Managed Service for Prometheus?

With Amazon Managed Service for Prometheus, there are no upfront fees or commitments. You pay only for what you use, based on metrics ingested, queried, stored, and collected. You will be charged at the end of the month for your usage.

Amazon Managed Service for Prometheus is Generally Available. We are charging for usage of metrics collected, ingested, stored, and queried. Customers will be charged according to the pricing table below.

AWS Pricing Calculator

Calculate your Amazon Managed Service for Prometheus cost in a single estimate.

Create your custom estimate now

Pricing

Amazon Managed Service for Prometheus counts each **metric** sample ingested to the secured Prometheus-compatible endpoint. Amazon Managed Service for Prometheus also calculates the **stored** metric samples and metric metadata in gigabytes (GB), where 1GB is 2^{30} bytes. The storage charge is determined by the Prometheus metrics samples (typically 1 or 2 bytes) and metadata. The Prometheus metric metadata varies in size based on the name of the metric and its associated labels (key/value pairs). Charges for metrics ingested and stored are prorated by the hour and billed only when you send metrics to Amazon Managed Service for Prometheus. Note that there is no Data Transfer IN charge for AMP.

Amazon Managed Service for Prometheus meters the Query Samples Processed (QSP) from all QueryMetric API requests. QSP is the total number of data points queried using Prometheus Query Language (PromQL) for a given time range.

For Amazon Managed Service for Prometheus collector, you are charged for the number of hours it is turned on and the number of samples collected.

The usage of Amazon Virtual Private Cloud (VPC) or VPC peering with Amazon Managed Service for Prometheus collector will incur additional charges as explained on the Amazon Elastic Compute Cloud (EC2) on-demand pricing page.

Region:

Europe (London)

Metrics Ingested	Cost ($/10M samples)
First 2 billion samples	$0.90
Next 250 billion samples	$0.35
Over 252 billion samples	$0.16

	Cost
Metric storage	$0.03/GB-Mo
Query Samples Processed (QSP)	$0.10/B samples processed
Collector hours	$0.04/collector hours
Samples collected	$0.03/10M samples

Except as otherwise noted, our prices are exclusive of applicable taxes and duties, including VAT, GST, and applicable sales tax.

AWS Free Tier

As part of the AWS Free Tier, you can get started with Amazon Managed Service for Prometheus for free (in supported AWS Regions). Free Tier customers receive:

Metric samples ingested	40 M
Query samples processed	200 B
Metrics stored	10 GB

Learn more about AWS Free Tier here.

Pricing examples

Note: Pricing values displayed here are meant to be examples only.

Example 1 - Amazon Elastic Kubernetes Service (EKS) on EC2 and Kubernetes

In this example you are monitoring 1 Kubernetes cluster with 10 nodes with 1000 Prometheus metrics per node collected every 30 seconds for the entire month (744 hours). This example does not factor in your AWS Free Tier benefits. The following usage and costs will be calculated for 1) metric sample ingestion, 2) storage, and 3) query.

Metric sample ingestion usage and costs

First, we will calculate the number of metric samples ingested and the costs for these samples. The monthly number of samples ingested is 892.8 million samples (10 nodes * 1000 metrics per node / 30 collection interval in seconds * 5,600 seconds in an hour * 744 hours in a month = 10 * 1000/30 * 5,600 * 744 = 892.8 million samples).

The monthly metrics ingested costs is $80.93 ($0.90/10,000,000 for the first 2 billion metric samples * 892.8 million samples).

Storage usage and costs

Second, we will calculate the storage in GB and the costs for the metric samples and metric metadata stored. For this example, we will assume that you have 20 labels per metrics with 100 unique label values per label. The metadata including the metric name, labels, and label values are stored daily. We will also assume an average of 30 bytes per label and label value and 2 bytes per metric sample ingested.

The monthly GB of storage is 0.25 GB (metric metadata * # of days per month) + (2 bytes * number of metric samples = (1000 metrics * 20 labels * 100 label values * 30 bytes * ~30 days per month) + (2 bytes * 892.8 million samples) = 1800 million bytes + 1,785.6 million bytes = 3,585.6 million bytes = 3.34 GB).

The storage charge is $0.10 ($0.03/GB * 3.34 GB).

Query usage and costs

We will now calculate the query samples processed and query costs for this example. We will assume you have 1 end user monitoring a dashboard for an average of 2 hours per day refreshing it every 60 seconds with 20 chart widgets per dashboard (assuming 1 PromQL query per widget). We will assume for this example that the samples processed per query is 100,000.

First, start by calculating the count of queries per month which is 72,000 queries (1 end user * 20 charts * 2 hours of viewing a dashboard * 3,600 seconds per hour / 60 seconds per dashboard refresh * ~30 days in a month).

Second, we calculate the query samples processed per month which is 7.2 billion (72,000 queries * 100,000/query)

We can now calculate the query charge of $0.72 ($0.10/billion samples processed * 7.2 billion samples processed).

We can now calculate your **monthly costs** for this example:

- Metric ingestion costs: $80.93 ($0.90/10M * 892.8 million samples)
- Storage costs: $0.10 ($0.03/GB * 3.34 GB)
- Query costs: $0.72 ($0.10/billion samples processed * 7.2 billion samples processed)

Total Monthly charges = $80.93 (metric ingestion) + $0.10 (storage) + $0.72 (query) = $81.75

Example 2 - Collecting Prometheus metrics from Amazon Elastic Kubernetes Service (EKS) with Amazon Managed Service for Prometheus

In this example you are monitoring 1 Kubernetes cluster with 10 nodes with 1000 Prometheus metrics per node collected every 30 seconds by Amazon Managed Service for Prometheus collector for the entire month (744 hours). The following usage and costs will be calculated 1) metric sample ingestion, 2) storage, and 3) Amazon Managed Service for Prometheus collector samples collected and collector hours.

Metric sample ingestion usage and costs

First, we will calculate the number of metric samples ingested and the costs for these samples. The monthly number of samples ingested is 892.8 million samples (10 nodes * 1000 metrics per node / 30 collection interval in seconds * 5,600 seconds in an hour * 744 hours in a month = 10 * 1000/30 * 5,600 * 744 = 892.8 million samples).

The monthly metrics ingested costs is $80.93 ($0.90/10,000,000 for the first 2 billion metric samples * 892.8 million samples).

Storage usage and costs

Second, we will calculate the storage in GB and the costs for the metric samples and metric metadata stored. For this example, we will assume that you have 20 labels per metrics with 100 unique label values per label. The metadata including the metric name, labels, and label values are stored daily. We will also assume an average of 30 bytes per label and label value and 2 bytes per metric sample ingested.

The monthly GB of storage is 0.25 GB (metric metadata * # of days per month) + (2 bytes * number of metric samples = (1000 metrics * 20 labels * 100 label values * 30 bytes * ~30 days per month) + (2 bytes * 892.8 million samples) = 1800 million bytes + 1,785.6 million bytes = 3,585.6 million bytes = 3.34 GB).

The storage charge is $0.10 ($0.03/GB * 3.34 GB).

Samples collected and collector hour costs

Third, we will calculate the collector hour costs. The agentless collector costs $0.04 per collector hour and $0.03 per 10M samples collected, and we assume a single collector for a single EKS cluster. Assuming 744 hours in a month, the monthly cost is: 1 * $0.04 * 744 + $0.03 * 892.8 million samples / 10M = $32.44.

The agentless collector charge is $32.44.

We can now calculate your **monthly costs** for this example:

- Metric ingestion costs: $80.93 ($0.90/10M * 892.8 million samples)
- Storage costs: $0.10 ($0.03/GB * 3.34 GB)
- The agentless collector costs: $32.44 ($0.04 per collector-hour * 744 hours + $0.03 * 892.8 million samples / 10M)

Total costs: $113.47

Example 3 - Calculating Query Samples Processed (QSP) and costs

In this example, you are running a query to calculate the CPU of your 10,000 nodes in the last hour and an alerting rule running every minute to notify you when one of those nodes exceeds 80% CPU capacity. This example does not factor in your AWS Free Tier benefits. We will calculate the usage and costs of running the query and the alerting rule.

Query usage and cost

First, we calculate the number of samples a single query processes by identifying the number of metric time series and the provided time range. A metric time series is identified by its metric name and key-value pairs. In this example, each node has a single CPU metric resulting in a query that processes 10,000 metric time series for a 1 hour time range. Prometheus queries process one sample per resolution step, which is the minimum amount of time between any two data points. For example, we'll assume a step of 1 minute, making it 1 sample per minute in the time range. This is 10,000 time series x 1 hour x 1 sample/min x 60 min/hour = 600,000 samples/query.

Assuming you ran the query 1,000 times, the cost would be $0.06 ($0.10/billion samples processed * 0.0006 billion samples * 1,000 queries).

Alerting rule usage and cost

We calculate the number of samples the rule processes per evaluation by identifying the number of selected metric time series. Since the rule identifies one CPU metric time series per node, it processes 10,000 time series per evaluation. Alerting rules in Prometheus run as instant queries, which are queries that are evaluated at a single point in time. Therefore, this alerting rule evaluates 10,000 time series x 1 sample/time series = 10,000 samples. Since this alarm runs every minute, this alerting rule processes 10,000 samples x 60 min/hour x 730 hours/month = 438,000,000 samples/month.

The cost of this alerting rule is $0.04 per month ($0.10/billion samples processed * 0.438 billion samples).

Get Started with AWS

Learn how to start using AWS in minutes

AWS Free Tier

Gain free, hands-on experience with AWS for 12 months

Proven Security to Accelerate Your Innovation

Find the resources you need to build, run, and scale your applications on the most secure cloud computing environment available today »

How to get started

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.