Understanding the Hidden Costs of Monitoring Cloud-Native Environments
A Costs Comparison Guide

sysdig

Table of Contents

01

Intro

Kubernetes and cloud services adoption has been increasingly accelerating in the last few years. Many businesses started to rapidly expand their cloud-native applications deployments, moving away from their old monolithic architectures. That brought tons of benefits, like huge improvements on the time to market, which let companies deliver their products at a much faster pace. Along with that drastic change in software-patterns-design-and-architecture mindset, a new challenge came out of nowhere: Traditional observability platforms are not valid anymore.

Observing cloud-native applications, Kubernetes services, or any of the cloud services available nowadays can be tough. Kubernetes is designed to rapidly scale up cloud-native applications, and in most cases, rely on the scheduler's artificial intelligence to dynamically decide the best place to run our cloud-native applications. Containers' lifespan is variable and shorter than monolithic applications by design. While containers can last only a few seconds by nature, monolithic applications were designed to run without interruption for years. That requires new and modern observability systems; instead of monitoring static monolithic and defined up-front architectures, those have to be able to not only capture container insights, but dig into the Kubernetes and cloud space, plus collect and process business metrics.

All that glitters isn't gold, though. New observability systems helps customers with monitoring efficiently cloud-native applications, but there is a new menace: Costs. This is a key factor when choosing a new observability platform. There are lots of cloud-native observability products out there, but not all of them are cost-effective enough. Complex pricing models and hidden cost policies are some of the factors that make it even harder to understand and forecast an observability bill. Quite often, businesses are not aware of this huge cost problem until the bill is on their table.

02

6 Monitoring Challenges Result in Higher Complexity and Costs

From an operational standpoint, there are a few common problems or situations that most businesses have to deal with during their observability journey. Let's talk about six observability challenges you may face while monitoring your cloud-native applications, or watching your own custom metrics.

1. Complexity of managing multiple tools for multiple clouds

Thanks to fast cloud adoption, cloud-native architectures are becoming increasingly complex compared to traditional designs. Many businesses opt to spread microservices across different cloud providers, or adopt hybrid models, which significantly reduces the risk against total outages, improving availability and reliability of their services.

In terms of observability, multi-cloud and hybrid strategies add complexity to the whole picture. Companies that opt for multi-cloud architectures need an observability platform to not only get metrics and monitor applications, but to obtain insights from their own cloud provider instances and services.

2. Scaling to monitor a growing environment is hard

Companies starting their monitoring journey with DIY Prometheus soon realize how complex it is to scale up their monitoring platform. In order to preserve optimal performance and keep up with the fast observability adoption within your organization, Prometheus may need to scale. Otherwise, depending on your observability service sizing, your Kubernetes and cloud-native applications metrics won't be available.

Prometheus doesn't scale well when the metrics volume grows. There are a few challenges that DIY Prometheus users need to address when managing large Prometheus metrics volumes. If you want to learn more about how to overcome these challenges, check out this Prometheus monitoring guide.

3. Finding the metrics you need most

DIY Prometheus users usually rely on the Prometheus service discovery mechanisms to start scraping metrics right away. Apart from that, they can also set up their own Prometheus configuration to pull metrics from custom endpoints. Prometheus exporters let practitioners start scraping metrics from third-party applications like a MySQL instance or from Kubernetes internal components like node-exporter or KSM. This approach is pretty straightforward, but what happens if any of the exporters you deployed suddenly stopped working? What about being confused between the five different exporters you found for the same application? Is that new exporter you just deployed secure enough? Does your observability service have you covered with Kernel and Kubernetes insights?

Prometheus exporters usually default to pull all the metrics available from its selected endpoints. That does not necessarily mean it is the best approach since most of the metrics add unnecessary noise and lead to useless data. In addition, the more data the the more risk of facing a cardinality explosion.

4. Beware of metrics cardinality explosion

Metrics are identified by a unique name and contain one or more labels, which may contain one or more values. The more labels in a metric and the more values for every label, the more cardinality. Metrics volume tends to grow over time, application usage intensifies, and developers tend to make label additions to improve traceability. The problem comes up when cardinality growth goes unnoticed. Metric changes like new label additions can burst cardinality, creating more and more new metrics in a short period of time. That's what we know as a cardinality explosion or a cardinality spike.

Cardinality spikes can jeopardize the stability of your monitoring platform, and even cause an outage if metrics volumes go wild. Cost is another huge problem with cardinality explosion; most SaaS observability services charge for metrics volume and metrics ingestion, skyrocketing the observability bill.

5. OSS is free but there are other costs around it

Open Source Software (OSS) is free; it's one of its key benefits, and one of the main reasons behind its popularity. Prometheus has become the de-facto standard for cloud-native applications observability in Kubernetes. You can easily deploy your own Prometheus instance and start monitoring your applications right away. Starting this monitoring journey is straightforward, and you don't need to pay for it.

When running Do It Yourself (DIY) Prometheus, there are other costs that need to be considered, however.

It doesn't matter whether the foundation of your business infrastructure is on-prem or you rely on cloud instances. You need to expense these infrastructure costs and reserve that cash from your budget. OSS is free but it needs to run somewhere and people to help run it, and this is not free of charge at all.

In terms of long-term storage costs, companies need data retention. It's key to look back on your metrics and compare trends one year back. Companies leveraging DIY Prometheus have to pay for their own storage to back up their metrics, plus the operational costs of maintaining it.

6. High workload and skills required for operations teams

As mentioned above, observability platforms tend to grow as time goes by. Scaling cloud-native applications and Kubernetes clusters, means more containers, applications, and services to monitor. In short, that's more metrics to manage. In order to make all this work, companies need engineers to manage, maintain, and operate the whole ecosystem.

Teams operating and maintaining the observability service incur costs as well. A group of humans are responsible for stability and availability of this service. That means hours and days of work. Those have to be paid, especially when on-call is required. Failing to provide enough engineers to maintain and operate the observability platform could end up in serious problems, such as performance issues and outages.

03

Why Your Monitoring Bill Keeps Rising

It's time to talk about some of the cost challenges you may find when using an observability solution. There are a lot of monitoring solutions available in the market, so some of these challenges may fit better or worse for your use case, depending on your observability platform and how you use it.

Beware of the observability hidden costs

When it comes to analyzing costs for an observability service, it is quite common to come across tons of different features, assets, or inherent functionality like query sample processing (QSP). This big catalog of assets, and the lack of understanding on which of these components may be charged, may incur a huge unexpected spending in your observability bill.

Some vendors apply charges for basic functionalities, like the already mentioned QSP by user, the total amount of long-term storage space, or even per container being monitored. These very basic assets like QSP are quite volatile, and they tend to fluctuate and grow over time. The more users and dashboards you manage, the more queries will be processed. Similarly, the more data retention you need, the more storage you'll need to pay. These extra costs can make your bill explode.

How much does your managed service for Prometheus cost?

Companies will find a lot of managed services for Prometheus available in the market; there is a wide variety of platforms, features, and prices. One of the biggest problems with Prometheus-managed services is the metrics ingestion interval time. It varies between providers. For example, Grafana Labs ingest metrics every 60 seconds, while others like Sysdig implement 10 seconds out of the box. Your total amount of time series will be defined by this metric interval period. The lower the interval, the more frequency on metrics ingestion, so the more data points you will collect. Some vendors, instead of reducing the price per time series (TS) to prevent large bills, prefer to increase the default interval time. The downside is you have less data.

On the other hand, many managed services for Prometheus providers charge for functionalities that are inherent to the product, like the QSP. The more users querying data, watching metrics via dashboards, and setting alerts, the larger the bill will be.

It's worth mentioning prices vary a lot between providers, as you'll see in the next section of this document.

Did your observability bill explode?

Companies may manage millions of custom metrics on a daily basis. Custom metrics are key for them since most of their business KPIs are measured that way. As custom metrics become more and more popular within the company, as we mentioned in the previous section, cardinality starts to grow. How observability platforms charge for custom metrics is very important. Cost-effective observability platforms will be able to contain companies' observability spending, while others may incur huge costs like this Datadog's customer, which had to face a $65M bill.

04

Managed service for Prometheus

In previous sections, we talked about the challenges you may face when managing and operating an observability platform for your Prometheus metrics. This time, go into detail on some of the managed services for Prometheus available in the market that can help you overcome all these challenges.

What is a managed service for Prometheus?

First of all, it's worth clarifying what you should expect from a Managed service for Prometheus, and how it differs from a DIY Prometheus platform. There are a lot of variants and offerings in the market, each providing its own features and capabilities. As you'll see later, prices vary among vendors. For now, let's summarize in general terms what you can expect from such a service.

Let's see it in detail.

- It's a Prometheus-compatible managed service. It does not necessarily have Prometheus software running on the cloud, but is a Prometheus-compatible ecosystem maintained and operated by a vendor.
- It can ingest your metrics from the available endpoints in your infrastructure. While most vendors allow either pull metrics via agent, or push metrics straightaway with remote write, some - like AWS - only support remote write.
- It provides long-term storage for your metrics. That way, you can inspect historical data whenever you need it.
- It allows you to watch and inspect your metrics, either running your own queries or visualizing data in dashboards.
- You can set alerts for your metrics based on your preferences and thresholds.

Sysdig's managed service for Prometheus pricing in detail

Sysdig Monitor charges metric ingestion per TS. Below, you'll find Sysdig's metrics pricing.

Component	Cost
Node exporter, cAdvisor, and KSM metrics included	$0
First 2K TS included in agent subscription per node	$0
> 2K TS per agent	$5 ($/1000 TS)
Sysdig Agent	$30/node

Note that **node exporter, KSM, and cAdvisor** time series metrics are **included in the Sysdig Agent price**. Hence, these **time series are not charged**. In addition, alerting and QSP are included, so you won't be charged for these features either.

Competitors and their pricing

Disclaimer: Here, you'll find the prices that correspond to the time this guide was written (June 2023). For current pricing, please check the public pricing information from every vendor.

AWS

The current Amazon managed service for Prometheus prices are available here. Amazon also provides its own pricing calculator to estimate your bill.

Amazon charges for different services within its Amazon managed service for Prometheus, like the metric ingestion costs, storage, and query. Metric ingestion is charged per sample.

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

Other costs	
Metrics storage	$0.03/Gb-Mo
Query Samples Processed (QSP)	$0.10/B samples processed

Google Cloud

Google Cloud monitoring pricing is available on its website. Google also has its own pricing calculator, just select "Cloud operations (Logging, Monitoring, Trace, Managed Prometheus)" and estimate your bill.

Google Cloud charges Metric ingestion per sample. The following table shows the Google Cloud monitoring pricing.

Metrics ingestion	Cost ($/1M samples)
First 50 billion (B) samples	$0.15
Next 50B-250B samples	$0.12
Next 250B-500B samples	$0.09
> 500B samples	$0.06

Some other items, such as "Monitoring API calls" usage and "Execution of monitoring uptime checks," may or may not be charged, depending on the usage within a full month.

Azure

Azure Monitor pricing is available here. If you want to estimate your Azure Monitor bill, you can use its pricing calculator. Azure charges metrics ingestion per sample.

Metrics ingestion	Cost ($/10M samples)
Any number of samples	$0.16

Queries	Cost ($/1B samples)
Metrics queries	$0.10

Azure charges for alerts and notifications as well –like emails, push notifications, or web hooks, among others.

Grafana Labs

If you want to check the Grafana Cloud pricing model, visit its website. On the pricing page, you can also calculate the estimated cost of your bill. Grafana Labs charges metrics ingestion per TS.

Metrics ingestion	Cost ($/1000 TS)
First 20K TS	Included in monthly usage subscription
Next 1K TS	$8

Grafana Labs charges based on active series, which may cause your bill to vary depending on your metric usage. On the other hand, alerting and QSP are included, so no extra costs are incurred.

Per Grafana Labs documentation, a time series is considered active if new data points have been received within the last 15 or 30 minutes.

Managed service for Prometheus cost comparison

K8s single cluster use case

Note that the number of time series you can have in your Prometheus instance can vary significantly and is dependent on your architecture. The more applications and operational tasks –like redeploy, creation, deletion, and scaling in your cluster, the more time series you will generate. Depending on how volatile your Pods and Kubernetes objects are, cardinality explosions may occur and can cause serious trouble.

For this use case, we have the following information about the Kubernetes infrastructure; this architecture will serve as a foundation to estimate the managed service for Prometheus costs for every vendor:

- One Kubernetes cluster
- 25 nodes

Next, you'll find the total number of time series registered in the Prometheus instance after a few days. This Kubernetes cluster was running under a normal load, not being stressed by heavy workloads or peaks of user activity. To emulate a minimal and a likely application lifecycle, we redeployed a few applications every day. These are the number of time series generated by job:

- kubernetes-apiservers: 73,713 TS
- kubernetes-pods: 275,421 TS
- kubernetes-nodes-cadvisor: 257,649 TS
- kubernetes-service-endpoints (node exporter + KSM): 144,202 TS
- kubernetes-nodes: 42,166 TS
- kube-dns: 370 TS
- etcd: 4399 TS
- prometheus: 1068 TS
- felix_metrics: 4008 TS
- kube_controller_metrics: 63 TS
- **Time series TOTAL: 803,059**

In terms of query processing and user activity, let's start with the following assumptions:

- 10 different users accessing Prometheus with their own dashboards and graphs reporting data for their projects.
- An average of eight graphs querying data per user.
- The refresh interval for each graph is 10 seconds. Assuming an average of two hours per user and day, it corresponds to 720 queries. 5,760 is the total number of queries for eight graphs per day and user.
- Data is being shown in a three hour timeframe on average.
- The highest number of samples processed in this environment for three hours is 867,303,720.
- We'll assume 300,000 samples on average per query.
- Monitoring query processing ~525,657,600,000[1]
- Alerting query processing ~5,256,000,000,000[2]

Once we have all the data, let's do some math!

First, we need to calculate the number of samples ingested based on our numbers.

803,059 TS / 10 collection interval in seconds * 3600 seconds * 744 hours in a month = 215,091,322,560 -> ~**215 billion samples**.

For those services where storage costs are charged, let's assume that storage initially needed for that volume of metrics is ~**12GB**.

[1] Monitoring query processing has been calculated from: 10 monitoring users * 5,760 queries per day and user * 30 days a month * 300,000 avg samples per query (3h).

[2] Alerting query processing has been calculated from: 2 executions per minute * 200 alerts * 60 minutes per hour * 730 hours per month * 300,000 avg samples per alerting rule.

With regards to queries processed, based on previous calculations, we'll assume that the total volume of queries is ~**5,781 billion queries / month**. You may think this number is too large but, on the contrary, it may be too small if we take into account that you may be querying millions of samples for a single query.

Notice that a 10-second sampling interval was used to calculate the total number of samples. Some vendors, like Grafana Labs, implement a 60-second sampling interval by default.

Disclaimer: Bellow, you'll find a quick calculation of the managed service for Prometheus costs for every vendor. Keep in mind that this is an approximation, and cost may fluctuate depending on the usage of your monitoring platform.

AWS

Let's see what the costs charged by each service would be.

Service	Cost
TS first 2B samples	2B samples * $0.90 /10M = $180
TS next 250B samples	213,091,322,560 samples * $0.35/10M = $7,458.19
Storage	$0.03 * 12.92GB * 365 days = $141.52 / month
Query Samples Processed (QSP)	$0.10/ B * 5,781 B queries = $578.1 / month
Total cost	**$180 + $7,458.19 + $141.52 + $578.1 = $8,357.81 / month**

So, if you own an Amazon managed service for Prometheus for processing the monitoring data belonging to the architecture defined earlier, you would spend around $8,357 a month.

GCP

It's now time to analyze the pricing for Google Cloud monitoring.

Service	Cost
TS first 50B samples	50B samples * $0.15 /1M = $7,500
Next 50B-250B samples	165,091,322,560 samples * $0.12/1M = $19,810.95
Total cost / month	**$7,500 + $19,810.95 = $27,310.95**

If you own a Google Cloud monitoring instance for processing the same data, you'll pay around $27,310 a month.

Azure

Let's analyze the Azure offering for monitoring and ingesting your Prometheus metrics.

Service	Cost
TS	215,091,322,560 * $0.16 /10M = $3,441.46
Query processing	$0.10/ B * 5,781 B queries = $578.1
Total cost	**$3,441.46 + $578.1 = $4,019.56 / month**

Using Azure for this specific use case would cost around $4,019 a month, and you'd also need to take into account the **costs related to alerts and notifications**, which are **extra assets that would be charged**.

Grafana Labs

These are the Grafana Labs costs for its Grafana Cloud product with a 10-second sampling interval.

Since the cost will vary depending on your active metrics, let's suppose all of your metrics (100%) are active.

Service	Cost
Service fee	$299 / month
TS first 20k	Included in subscription
TS > 20k	803,059-20k = 783,059 * $8 / 1,000 = $6,264.4
Total cost / month	**$299 + $6,264.4 = $6,563.4**

With Grafana Labs, total cost would be around $6,563 a month.

Sysdig Monitor

Sysdig Monitor implements a 10-second sampling interval by default, resulting in up to 6x more metrics compared to competitors and a lower cost, as you'll see next.

First of all, you need to pull out node exporter, cAdvisor, and KSM metrics from the current numbers —Sysdig Monitor brings its own custom equivalent metrics. For the sake of simplicity, let's subtract the following jobs from the total number of time series:

- kubernetes-nodes-cadvisor: 257,649 TS
- kubernetes-service-endpoints (node exporter + KSM): 144,202 TS
- kubernetes-nodes: 42,166 TS

The **new TOTAL number of time series** is: 803,059 - 257,649 - 144,202 - 42,166 = **359,042 TS**. The number of billable time series has been **reduced by ~55%**!

Service	Cost
Agent cost	$30 * 25 nodes = $750
Metrics included in agent subscription	2,000 * 25 nodes = 50,000 TS included free of charge
Metrics ingestion	Next 309,042 * $5 / 1,000 = $1,545.21 / month
Total cost	**$750 + 1,545.21 = $2,295.21 / month**

If you use Sysdig Monitor as your managed service for Prometheus, you'd pay around $2,295 a month.

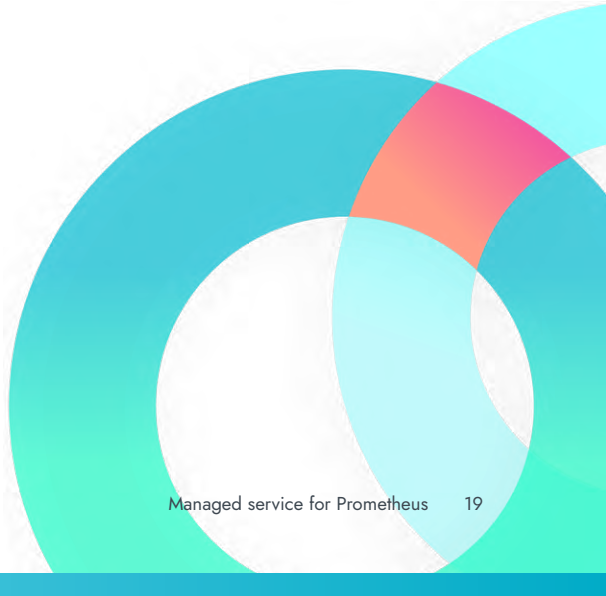

Price comparison

After doing all the calculations, it's time to sum up the costs of every service. This time, the price is calculated by time series for different ingestion intervals (60s, 30s, and 10s).

TS calculator	AWS First ~44,800/ 22,400/ 7,500 TS	AWS Next ~933,393 TS	GCP First ~1,120,071/ 560,035/ 186,678 TS	GCP Next ~746,714 TS	Azure - any number of TS	Grafana Labs & Sysdig	QSP	Disk	TOTAL
AWS (60s)	$180.00	$1,184.70					$578.17	$141.52	$2,084.38
AWS (30s)	$180.00	$2,439.40					$578.17	$141.52	$3,339.08
AWS (10s)	$180.00	$7,458.20					$578.17	$141.52	$8,357.88
GCP (60s)			$5,377.28						$5,377.28
GCP (30s)			$7,500.00	$2,603.65					$10,103.65
GCP (10s)			$7,500.00	$19,810.96					$27,310.96
Azure (60s)					$573.58		$578.17		$1,151.74
Azure (30s)					$1,147.15		$578.17		$1,725.32
Azure (10s)					$3,441.46		$578.17		$4,019.63
Grafana Labs (60s)						$6,563.47			$6,563.47
Sysdig (10s)						$2,295.21			$2,295.21

When comparing managed service for Prometheus costs, be aware of the costs by metric interval sampling. Sysdig's metric interval sampling is **10 seconds by default**, while DIY Prometheus, GCP, and Grafana pull metrics every 60 seconds. Despite collecting 6x more data than some of its competitors, Sysdig is a much more cost-effective option. There are no extra charges for storage or queries samples processing with Sysdig. Bear in mind that **these services can dramatically increase your bill**, and with the inherent complexity of forecasting QSP numbers, it is something variable that depends on the users' usage.

Comparing all the managed service for Prometheus prices analyzed in this guide under the 10-second metrics ingestion scope, you'll see a huge difference among vendors:

- AWS: ~$8,357 / month
- GCP: ~$27,310 / month
- Azure: ~$4,019 / month
- Grafana Labs: ~$6,563 / month
- Sysdig: ~$2,295 / month

Azure is almost 2x more costly than Sysdig, AWS is almost 4x, and GCP is 12x more.

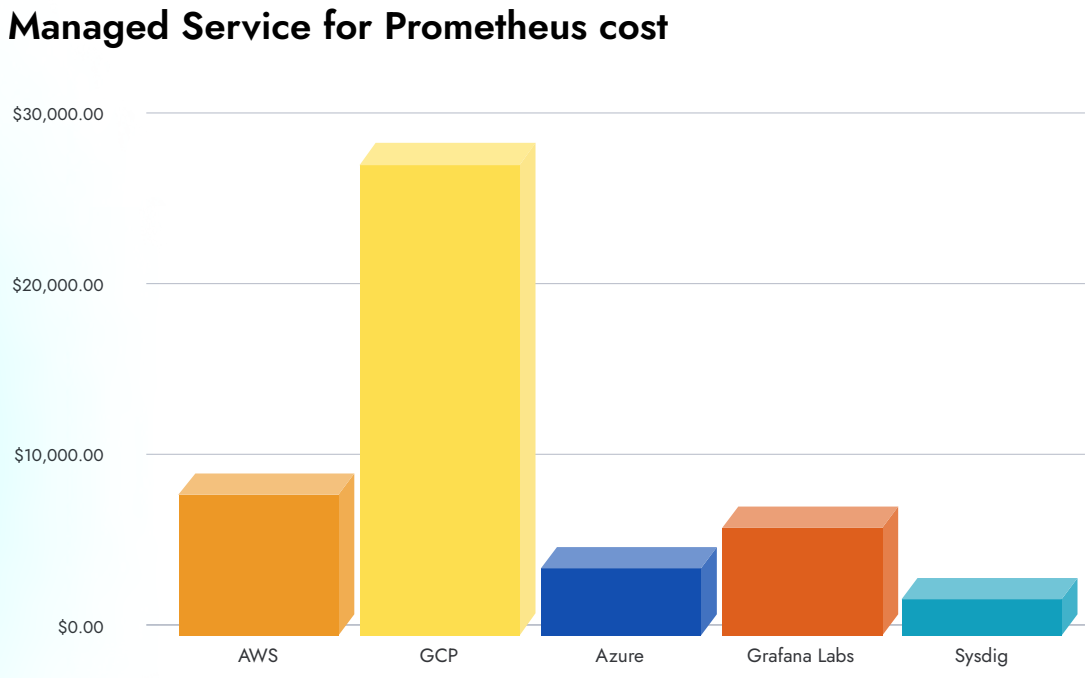

K8s multi-cluster use case

In the previous section, you saw how time series are generated and how volatile and dynamic the numbers can be for every use case. This time, we'll analyze costs in a larger scenario. This architecture is made up of five K8s clusters, 50 nodes in each cluster.

- Five Kubernetes clusters
- 50 nodes per cluster

The total number of time series have been calculated under regular load, no stress or high load average peaks. For the sake of simplicity, during the testing cycle we redeployed/scaled down/scaled up a few deployments. That way, we can emulate a real application lifecycle. Bellow, you'll find the number of time series generated by this group of K8s clusters by job:

- kubernetes-apiservers: 368,565 TS
- kubernetes-pods: 4,131,315 TS
- kubernetes-nodes-cadvisor: 3,607,086 TS
- kubernetes-service-endpoints (node exporter + KSM): 2,018,828 TS
- kubernetes-nodes: 505,992 TS
- kube-dns: 1,850 TS
- etcd: 21,995 TS
- prometheus: 5,340 TS
- felix_metrics: 20,040 TS
- kube_controller_metrics: 315 TS
- **Time series TOTAL: 10,681,326**

In terms of query processing and user activity, let's start with the following assumptions:

- 50 different users accessing Prometheus with their own dashboards and graphs reporting data for their projects.
- An average of eight graphs querying data per user.
- The refresh interval for each graph is 10 seconds. Assuming an average of two hours per user and day, it corresponds to 720 queries. 5,760 is the total number of queries per day and user.
- Data is being shown in a three hour timeframe on average.
- The highest number of samples processed in this environment for three hours is 11,535,832,080.
- We'll assume 300,000 samples on average per query.
- Monitoring query processing ~2,628,288,000,000[3]
- Alerting query processing ~26,280,000,000,000[4]

[3] Monitoring query processing has been calculated from: 50 monitoring users * 5,760 queries per day and user * 30 days a month * 300,000 avg samples per query (3h).

[4] Alerting query processing has been calculated from: 2 executions per minute * 200 alerts per cluster * 60 minutes per hour * 730 hours per month * 300,000 avg samples per alerting rule.

Disclaimer: Bellow, you'll find a quick calculation of the managed service for Prometheus costs for every vendor. Keep in mind that this is an approximation, and cost may fluctuate depending on the usage of your monitoring platform.

This time, we'll go straight to the point and calculate the costs for each provider using the price per time series we got from the previous section.

Price comparison

Let's compare the 10-second metrics ingestion prices for every service.

TS calculator	AWS First ~44,800/ 22,400/7,500 TS	AWS Next ~5,600,358/ 2,800,179/933,393 TS	AWS over 252B samples	GCP First ~1,120,071/ 560,035/186,678 TS	GCP Next ~4,480,286/ 2,240,143/746,714 TS	GCP Next ~5,600,358/ 2,800,179/933,393 TS	GCP over 500B samples	Azure - any number of TS	Grafana Labs & Sysdig	QSP	Disk	TOTAL
AWS (60s)	$180.00	$8,750.00	$3,597.03							$2,890.83	$1,882.28	$17,300.13
AWS (30s)	$180.00	$8,750.00	$11,226.06							$2,890.83	$1,882.28	$24,929.16
AWS (10s)	$180.00	$8,750.00	$41,742.18							$2,890.83	$1,882.28	$55,445.29
GCP (60s)				$7,500.00	$24,000.00	$20,413.30	$0.00					$51,913.30
GCP (30s)				$7,500.00	$24,000.00	$22,500.00	$27,217.73					$81,217.73
GCP (10s)				$7,500.00	$24,000.00	$22,500.00	$141,653.18					$195,653.18
Azure (60s)								$7,629.03		$2,890.83		$10,519.86
Azure (30s)								$15,258.06		$2,890.83		$18,148.89
Azure (10s)								$45,774.18		$2,890.83		$48,665.01
Grafana Labs (60s)									$85,589.61			$85,589.61
Sysdig (10s)									$27,747.10			$27,747.10

In summary, these are the costs for every service ingesting ~11 million metrics with a 10-second interval.

- AWS: ~$55,445/month
- GCP: ~$195,653/month
- Azure: ~$48,665/month
- Grafana Labs: ~$85,589/month
- Sysdig: ~$27,747/month

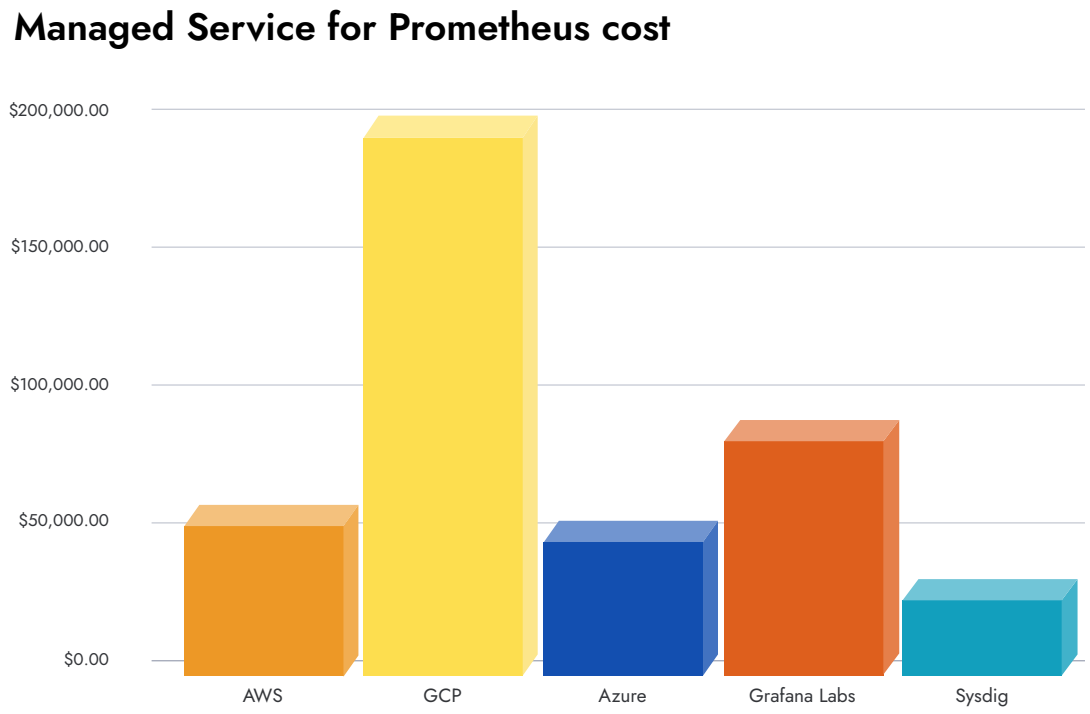

For a ~11 million time series volume, Sysdig's managed service for Prometheus is significantly cheaper than its competitors. GCP is almost 4x more expensive than AWS and Azure. Grafana Labs increased the cost significantly, being the second most expensive option. In comparison, Sysdig offers the most cost-effective solution.

Sysdig is Much More than a Managed Service for Prometheus

As you can see in the analysis above, Sysdig is the most cost-effective solution for processing and managing your Prometheus metrics. But that's not all. Sysdig Monitor brings some extra features which add an immense value to its customers. Let's see what we are talking about.

- While some service providers like AWS only allow users to ingest metrics via remote write, Sysdig is able to either receive metrics from remote write sources, or automatically pull metrics by itself.
- Sysdig Monitor has the shortest metrics ingestion interval time; it ingests metrics every 10 seconds out of the box. Others, like Grafana Labs, ingest data every 60 seconds. This means 6x less data than Sysdig.
- Sysdig Agent leverages eBPF for collecting metrics and getting kernel insights. Thanks to eBPF implementation, security and performance is improved, avoiding unnecessary risks when pulling kernel internal data.
- Tons of out-of-the-box dashboards and alerts for your infrastructure, cloud provider, third-party software, applications, and custom metrics.
- PromQL UI for those users that want to explore their data on their own, and a form-based UI for beginners that want to watch their data and learn PromQL at the same time.
- Sysdig Advisor is a troubleshooting tool that helps users with analyzing the root cause of their problems. It also provides recommendations and guidance on how to fix issues. It is included at no extra cost.
- Sysdig Monitor allows users to capture container insights every time an alert is triggered. DevOps teams can inspect this wireshark-like capture to get a better understanding of the problem thanks to the fine-grain data (syscalls).
- Sysdig users can monitor their Kubernetes and Cloud spending with the help of Cost Advisor[5]. This tool analyzes your historical data along with the cost information from your cloud provider and provides custom cost analysis and remediation steps for rightsizing your workloads.
- Observe all your Cloud and Kubernetes applications, services, and workloads from a single pane of glass. Out-of-the-box integrations are available for Sysdig Monitor customers.
- Sysdig Monitor is a multi-cloud observability platform. Connect, pull, and visualize cloud metrics and data with the out-of-the-box dashboards from the three main cloud providers (AWS, Google Cloud, Azure).

[5] Sysdig users interested in Cost Advisor can get access to it by purchasing a separate SKU. Contact your Sysdig sales representative for more information.

05

Custom Metrics Observability

Some businesses nowadays opt for other non Prometheus-compatible solutions, like Datadog. While Datadog is not a managed service for Prometheus, it offers tons of features that may go unnoticed for users who may never use them. On the other hand, some of the most important functionalities for many customers, like custom metrics, are sold at a very high price.

If you feel you are paying too much for your custom metrics, keep reading and discover how to reduce your custom metrics bill.

What are custom metrics and why companies need them

Custom metrics are a key component for many companies. Stock available in warehouses, shopping cart status, number of products sold, and operational status for industrial machines are some of the many KPIs that companies need for their own business tracking purposes.

As custom metrics become more and more popular within a company, developers start to instrument more applications, or add new labels to their existing metrics. This makes metric cardinality grow exponentially, and as you already read earlier in this document, this expected or sometimes unexpected, unnoticed growth may cause a lot of harm.

For many companies, custom metrics mean the largest asset in their observability bill. Managing and processing custom metrics with a cost-effective solution may help businesses to reduce their custom metrics bill by up to 75%.

Now, let's take a look at the price comparison between two observability platforms widely known for managing and processing custom metrics: Datadog and Sysdig.

Custom metrics pricing in detail

Disclaimer: Here you'll find the prices that correspond to the time this guide was published (June 2023). For current pricing, please check the public pricing information from every vendor.

As you will see in the table below, Datadog charges for different assets, like the product license, custom metrics store (beyond 200 custom metrics / host), ingestion, and container monitoring. Sysdig only charges for its license and the custom metrics beyond the free bucket. For further information on actual Datadog pricing, check here.

Let's see how these charges are spread.

Component	Datadog	Sysdig
Product license	$27/host	$30/host
Custom metrics	$5/100 custom metrics	$5/1000 custom metrics
Custom metrics ingestion	$0.10/100 ingested custom metrics	Free, unlimited at no extra cost
Free custom metrics bucket	200 custom metrics / host	2000 custom metrics / host
Container monitoring	$1/per container	Free, unlimited at no extra cost
Free container monitoring bucket	10 containers / host	N/A

While Sysdig Monitor users that exceed these 2,000 custom metrics per host will be charged at $5 for 1,000 custom metrics ($0.005 per custom metric), Datadog charges $5 for only 100 custom metrics blocks – $0.05 per custom metric. In short, Datadog custom metrics pricing is 10x more expensive than Sysdig.

In terms of metrics coming from integrations, Datadog doesn't charge for those metrics pulled from their out-of-the-box integrations. Sysdig considers custom metrics all of the metrics pulled from your platform, but includes at no extra cost: cAdvisor, node exporter, and KSM equivalent metrics, plus Kernel insights from your platform for monitoring and troubleshooting. Out-of-the-box integrations metrics are curated to ensure customers only get what they need, reducing noise and operational burden.

Apart from custom metrics, in terms of costs, Sysdig doesn't charge for other features. Users only have to pay the price per agent, which includes other features like Advisor, out-of-the-box dashboards, alerts, metrics enrichment, integrations, etc. These functionalities are available from the very beginning at no extra cost.

Custom metrics cost comparison

Let's first take a look directly at custom metric costs, as many of our customers found custom metrics to be by far the largest component of their monitoring bill.

In this example, we compare costs for 10 million custom metrics from your applications with Datadog and Sysdig. That's all, just custom metrics prices. The other pricing components listed in the table above will need to be considered when monitoring your infrastructure with both Datadog and Sysdig. We'll see this in more detail in a later example.

Vendor	Custom metric price	Custom metrics volume	Custom metric cost / month	Custom metric cost / year
Datadog	$5 / 100 = $0.05	10,000,000	$500,000.00	$6,000,000.00
Sysdig	$5 / 1000 = $0.005	10,000,000	$50,000.00	$600,000.00

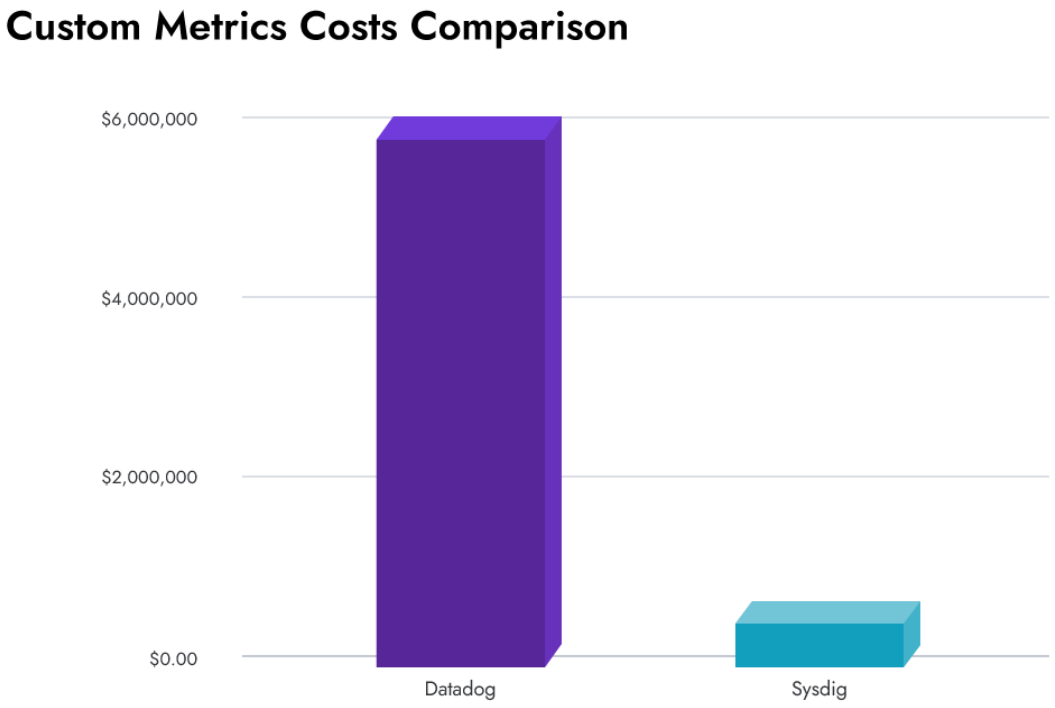

As anticipated in the previous section, Sysdig custom metric costs are 10x cheaper than its competitor. Businesses managing around 10 million time series would pay around $6,000,000 per year for custom metrics with Datadog, while Sysdig would reduce their spending 10x.

Check out this real customer use case and discover how Sysdig Monitor helps customers significantly lower their custom metrics costs while increasing their custom metrics volume.

Custom metrics cost comparison in detail

This time, we'll show you a custom metrics costs comparison in detail.

In this environment, there are five K8s clusters with the following architecture definition:

- Five Kubernetes clusters
- 50 nodes per cluster
- ~9,250 containers in total

Our observability platform is processing around 20M time series from applications running in these environments. There are many other components in these K8s clusters that generate Prometheus metrics as well, like the K8s control plane, cAdvisor, node exporter, KSM, coreDNS, etc. For the sake of simplicity, you'll find approximate numbers of metrics for every component in the following list.

- Application custom metrics: 20,000,000 TS
- K8s control plane: 530,000 TS
- K8s cAdvisor: 3,800,000 TS
- K8s Node exporter & KSM: 2,300,000 TS
- K8s nodes: 550,000 TS

Next, you'll find a price comparison between Datadog and Sysdig for monitoring a Kubernetes architecture like the one described in the heading of this section.

Vendor	Custom metrics cost	Ingestion cost	Containers	Host/ Agent	Total / month	Total / year
Datadog	$997,500.00	$19,959.00	$6,750.00	$6,750.00	$1,030,950.00	$12,371,400.00
Sysdig	$100,150.00	$0.00	$0.00	$7,500.00	$107,650.00	$1,291,800.00

Digging deeper into the numbers used for the math, as anticipated earlier, 20,000,000 TS is the total number of custom metrics in this environment for applications. For Datadog calculations, after applying 200 TS reduction per node (50,000 TS in total), it becomes **19,950,000 billable TS at $5/100TS, ~$997,500 for custom metrics**. Keep in mind that metrics ingestion, containers, and infrastructure hosts are also charged – check these extra costs in the table above.

As explained in the previous example, Sysdig Monitor charges a **much lower price**. That's **$5/1000TS** for custom metrics. Sysdig brings its own set of equivalent KSM, node exporter, and cAdvisor metrics at no extra cost, so these don't need to be ingested at all. In addition, as discussed in previous sections, the license includes 2,000 TS per node – that means 500,000 for free in this specific scenario.

Have you already noticed the differences in pricing between both products? While Datadog total costs are around **$1,030,950/month** and **$12,371,400** for a whole year, Sysdig charges **$100,150** and **$1,291,800** respectively for the same period. This is **10x cheaper** than the competitor.

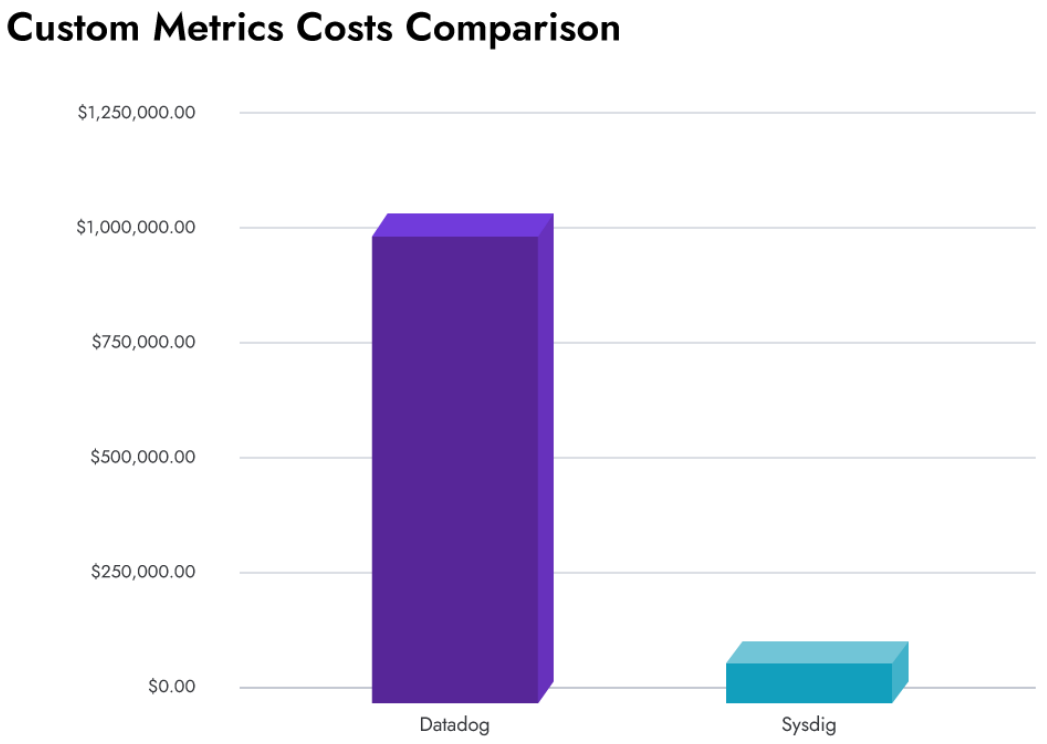

Sysdig's extra value for custom metrics

Sometimes, companies are not looking for a managed service for Prometheus but are using other observability platforms like Datadog. Those businesses managing a high volume of metrics with Datadog are exposed to paying very large bills. It's even worse for those companies that haven't faced a cardinality explosion yet; they'll have to face this cost explosion sooner or later. If you are in that position and looking for a cost-effective monitoring solution for your custom metrics, Sysdig can help with that and much more.

For the sake of simplicity, let's enumerate only a few points where Sysdig Monitor adds extra value. The summary done in the previous section, "What are the extra benefits Sysdig brings?" also applies here. Check it out first if you missed it.

- Control the metrics cardinality with the out-of-the-box dashboards and alerts, and avoid a huge wasted spending. Cardinality explosion won't be a problem anymore.
- Sysdig Monitor customers don't have to pay for long-term storage or QSP. If there is a cardinality explosion, you won't incur extra cost for storing your metrics.
- All the metrics, including your custom metrics, have Kubernetes and Cloud context. Thanks to this metric enrichment, Sysdig Monitor customers can troubleshoot issues and correlate Kubernetes and Cloud data with their metrics.

06

Conclusion

When it comes to controlling your observability platform costs, understanding how your observability provider charges per asset and functionality is key. Most of the problems that observability customers have are due to complex and hidden pricing policies, leading customers to jump out of a plane without a parachute.

As you saw in this guide, there are many observability solutions for your Prometheus metrics, but not all of them are cost-effective enough. Sysdig Monitor is not only a 100% Prometheus-compatible managed service, but a complete observability solution for any environment. Other than Prometheus and custom metrics observability, it also provides multiple functionalities to monitor and troubleshoot your Kubernetes, VMs, and Cloud providers.

If you are interested in learning more about how to cut custom metrics costs by up to 75%, we encourage you to watch this webinar.

Besides the cost efficiency while ingesting, processing, and managing Prometheus metrics, Sysdig allows Kubernetes and Cloud customers to monitor costs at infrastructure, cluster, namespace, and workload level. Discover your wasted spending and follow remediation steps to start reducing your Kubernetes and Cloud bill by 40% on average.

How Sysdig Monitor Helps

Sysdig provides a fully compatible Prometheus SaaS solution with long-term support which is permanently being improved to scale to millions of time series. It is also much more cost-effective than other Prometheus solutions available in the market. Sysdig Monitor helps you observe more and spend less. You can request a 30-days trial account and try it for free.

sysdig.com/start-free