

CLOUD MONITORING: CHOOSING THE RIGHT METRICS



[Cloud environments](#) are notorious for their lack of visibility and control. This makes it difficult to:

- Understand your cloud usage and strategy.
- Optimize for cost, performance, and [security](#).

Cloud computing data centers, and the services running on top of them, collect valuable logs of information that can help users identify and respond to infrastructure performance and changes in real-time, before the impact spreads across the network.

So, let's look at cloud monitoring and how to choose the right metrics for your organization.


What is cloud monitoring?

[Monitoring](#) is essential to maintain the health of the IT environments as well as the performance of apps and services operating on the cloud environment. Specifically, cloud monitoring technologies are designed to manage the [cloud infrastructure environment](#) and provide useful information on the performance of the hosted apps.



Cloud Infrastructure

Infrastructure components for cloud



for the following activities:

- SLA management
- Resource management, planning, and provisioning
- Datacenter management
- Troubleshooting
- Billing and accounting
- Security management
- Performance management

The cloud monitoring information is used

Metrics for cloud monitoring

A part of the decision making involves the practice of [collecting, analyzing, and understanding logs](#) across these metrics allows users to monitor cloud environments. Metrics provide knowledge on capabilities such as uptime, availability, quality of service, reliability and other key properties of the services being delivered over the Internet.

This information is collected at two levels:

- At the High Level, cloud monitoring assesses software resources on the virtual platform running in the cloud environment.
- At the Low Level, cloud monitoring assesses the underlying hardware infrastructure of the cloud computing datacenters.

Common cloud monitoring metrics

Some of the common metrics used in cloud monitoring are relevant in the following categories:

Metrics on the virtual machine

The performance of applications running on [cloud virtual machines \(VMs\)](#) depends on the underlying host performance. These host servers and hardware resources are shared across servers. Some VMs may consume these resources excessively or require additional resources.

In order to commit resources optimally and avoid resource bottlenecks on individual VMs, view the VMs collectively. The metrics should also support:

- Automated resource management
- Provisioning
- Autoscaling capabilities

For instance, remove alerts for downtime on specific VMs when the autoscale is designed to turn it off while it's not in use for prolonged periods.

Metrics from the cloud vendor

Service providers offer limited visibility and control into the underlying hardware of cloud systems. A variety of metrics information however is also available in the form of intuitive dashboards and reports. Users can choose metrics by types, categories, and groups that are generated frequently, time-stamped, and aggregated with the necessary descriptive information.

These details make it easy for users to interpret raw metrics and transform vast logs of information into insights through data analytics capabilities. Detailed analysis requires relevant skills and monitoring tools offered by cloud vendors.

Metrics on application performance

At the high level, metrics generate valuable information into [application performance](#). This information can be used to:

- Fine-tune applications
- Identify issues
- Escalate problem resolution

For example, it's important to capture transactions accurately in large-scale cloud environments. APM products may not capture this information for every user all of the time. Even when transactions are sampled at every small period of time, such as per-second, the number of transactions captured may only be a small percentage of millions of transactions that occur in a large-scale app every day.

Even when captured, additional analytics tools are required to analyze the complete information. The performance management capabilities of the [application performing monitoring \(APM\) tooling](#) can then help make insightful decisions.

Guidelines for custom cloud monitoring metrics

Vast volumes of metrics logs generated in large scale cloud environments are overwhelming and full of noise, which can cause unnecessary decisions that impact business performance. In the real world, your cloud monitoring strategy should guide the collection of useful metrics that solve key

business problems, help understand customers, and improve the user experience and business growth.

Use these guidelines when developing your cloud monitoring strategy:

1. Understand the cloud problem

The important first step is to understand the [implications of cloud migration](#), in terms of its differences from [infrastructure monitoring](#). Cloud computing is a fundamental shift in IT operations. Interactions with tools, metrics, functions, and the resource components can be short-lived but present significant implications to long-term business growth.

2. Measure what's important

Metrics help capture information of a system as it evolves. It's important to find metrics that collectively provide accurate, contextual, and insightful information on various aspects of cloud performance. Therefore, you must first understand the important attributes and parameters of the cloud performance and how they relate with your overall business performance and user experience.

3. Cut through the noise

Since cloud monitoring generates logs at overwhelming rates for analysis using conventional tools, advanced solutions for cloud monitoring are a business imperative.

- Technologies that help cut through the noise and avoid false alerts and unnecessary red flags are important to make the most out of your cloud investments.
- Advanced AI capabilities can help identify and understand patterns of performance, and adjust alerting mechanisms proactively based on real-time usage of the cloud infrastructure.

4. Understand performance on multiple levels

Investigate cloud infrastructure performance systematically. The cloud stack can be seen in terms of multiple layers:

- Infrastructure
- Platform
- Operating system
- Software layers

At each layer, identify the systems that perform the useful work and the underlying or interdependent systems that support that work. Capture metrics data that help investigate resource performance across these layers and resource components systematically.

5. Fix but don't forget

It's important to improve on a continuous basis:

- Incorporate feedback on current improvements into your future cloud management strategy.
- Correlate metrics data with previously captured information.

- Understand how the application stack and infrastructure environment has evolved as the user base scales.

If the symptoms of a current problem are recorded, the repeat of future problems demonstrating similar symptoms can be fixed efficiently.

Additional resources

For more on related topics, explore these resources:

- [BMC Multi-Cloud Blog](#)
- [State of IT Monitoring in 2020](#)
- [Observability vs Monitoring: What's The Difference?](#)
- [Choosing IT Metrics That Matter](#)