

Why Your Datadog Bill Keeps Climbing—and How to Stop It

Datadog can become very expensive when all hidden costs are added up. Here are the three main reasons why Datadog customers are transitioning away from Datadog:

1. Complex pricing and (almost certain) overage costs

Datadog's pricing structure can be complex, and users—especially those in large organizations or with distributed, cloud-native systems—will encounter unforeseen costs. One aspect of this complexity is the host based licensing for Infrastructure Monitoring and Application Performance Management products, where Datadog customers are billed for the number of hosts in their infrastructure stack and the number of hosts applications are deployed.

While this licensing model worked well in traditional architecture with monolithic applications, distributed cloud-native environments are highly dynamic. Auto scaling policies can increase or decrease host count elastically and on-demand.

Datadog bills for the maximum of 99 percent high-water mark of hosts that can explode your costs. Consider the following scenario: (example next page)

Scenario

Suppose you have varying host counts over a 30-day month.

Total Hours in Month

30 days × 24 hours = 720 hours

Hourly Host Counts

- **650 hours:** 50 hosts
- **60 hours:** 70 hosts
- **10 hours:** 90 hosts

Step-by-Step Calculation

1. Calculate Top 1% of Hours:

- 1% of 720 hours = 7.2 hours
- Round up to **8 hours** to cover at least 1%.

2. Exclude Top 1% (8 Hours):

- The highest host counts occur during the 10 hours with 90 hosts.
- Exclude the **8 hours** with the highest host counts (all from the 10 hours with 90 hosts).

3. Remaining Hours After Exclusion:

- **712 hours** (720 total hours - 8 excluded hours)

4. Determine the High-Water Mark:

- After exclusion, the highest host count is from the remaining 2 hours with 90 hosts and the 60 hours with 70 hosts.
- The maximum host count among these is **90 hosts** (from the 2 hours that weren't excluded).

5. Billable Host Count:

- **90 hosts**

You are on the hook for 90 hosts even though you used 90 hosts for just 10 hours in a month!

But, that's not all! Every host has limited entitlements such as the number of containers or span events. If you exceed those limits, you will be paying overages.

In contrast, Observe offers ingest based pricing that gives you a lot of flexibility. For example, in the case of metric data, the **average** of data points ingested per minute over one month.

2. Cost and Observability trade-offs

Even Datadog admits that no one can afford to ingest and retain all their data. Datadog recommends aggressive sampling. Better yet, it does that by default for you. 99% of your ingested spans are dropped by Datadog within 15 minutes as mentioned in their [pricing documentation](#):



Ingestion means sending your traces to Datadog and having all of them available for Live Search and Analytics for 15 minutes ... Retention means storing your most important traces (e.g. error, high latency, or business-critical ones) and making them available for search and analysis during a retention period of your choice (15 days by default)."

If you need to retain more trace data in Datadog, you can do so for an astronomical price:

- "7-day retention at \$1.27 per million spans per month (billed annually)"
- "15-day retention at \$1.70 per million spans per month (billed annually)"
- "30-day retention at \$2.50 per million spans per month (billed annually)"

Similarly, retaining log data for extended periods in Datadog can become prohibitively expensive, so Datadog recommends extracting metrics from logs for historical analysis.

But it [charges](#) every log derived metric as a custom metric. So you might reduce your log management bill but your custom metrics charges will skyrocket.

With Datadog you constantly need to make tradeoffs on comprehensive observability vs. costs. These tradeoffs result in suboptimal observability experience, blind spots, and higher MTTR.

Observe offers Zero-compromise Observability. Observe's modern cloud-native architecture enables faster data access and exploration, effortless scalability and dramatically lower cost. You don't need to sample or tier data avoiding blind spots. All data is always kept hot and can be retained for longer periods cost-efficiently.

3. Additional costs for premium support

Reliable technical support is crucial for business success. With Datadog, premium support is an add-on, [costing an additional 8% of your monthly spend](#). As your product usage and spending increase, so does the cost of obtaining quality support. This can be burdensome for large organizations or those with rapidly growing infrastructure, where support costs can escalate. You need to set aside an additional 8% of your monthly Datadog bill to get support!

Observe, on the other hand, provides premium support to guide you through your modernization and growth journey, often with no additional charge. From legacy tool migration to data onboarding and alert setting, you can count on Observe's customer support.

Loved by the Most Innovative Companies

Edgio

reveal



 **AUDITBOARD**



Capital One is focused on building seamless customer experiences that make banking and commerce simpler and easier—whether those customers are, digital or in-person. Like many cloud-first organizations, our data volume continues to expand. Observe provides a centralized and pre-correlated data layer that meaningfully organizes telemetry data from many sources at scale, helping drive faster response times.”

— Mark Cauwels, Managing Vice President, Enterprise Platforms Technology

