



Datadog Announces GPU Monitoring to Help Businesses Optimize Spend and Performance as They Aim to Scale AI Projects

April 22, 2026 at 9:00 AM EDT

The launch of Datadog's GPU Monitoring helps teams plan capacity, troubleshoot issues quickly, prevent costly failures and avoid wasted spend

NEW YORK, April 22, 2026 (GLOBE NEWSWIRE) -- [Datadog, Inc.](#), (NASDAQ: DDOG), the leading AI-powered observability and security platform, today announced that [GPU Monitoring](#) is available to customers everywhere. The new product addresses one of the most prevalent issues facing organizations today as they look for a scalable and effective way to manage expanding AI costs.

"GPU instances account for [14 percent of compute costs](#)—which is a huge issue as companies are struggling to build AI-first technology in scalable and smart ways. While these companies can see their costs climbing, they can't chargeback GPU spend across business units, see workload context or identify clear next steps for improvement. As a result, it is very challenging to budget and plan in thoughtful ways," said Yanbing Li, Chief Product Officer at Datadog.

The launch of GPU Monitoring marks one of the first times a single solution provides unified visibility across the AI stack—giving customers a single view linking GPU fleet health, cost, and performance directly to the teams relying on them for faster troubleshooting of slow workloads and cost savings.

"Smartly managing AI spend becomes a board-level conversation when capacity is misallocated, training and inference workloads stall, and costs escalate. We all know managing GPU costs is a huge problem we need to solve, but most companies are experimenting with solutions and it is still very difficult to get a single view of what is happening across the stack. GPU Monitoring fixes that with efficiency and reliability that we haven't seen before," said Li.

Today, most GPU tools provide high-level device health metrics, but they don't surface cross-functional resource contention issues, explain why training and inference workloads fail, or provide visibility into which devices are idle or ineffectively used. This lack of visibility slows down investigations and means that teams overprovision as the safest default—leading to wasted spend.

GPU Monitoring streamlines this work by linking fleet telemetry directly to the workloads consuming those resources, and gives platform engineering and machine learning teams a shared view to investigate together, enabling them to:

- **Scale AI without overspending:** With visibility and forecasting based on the usage patterns of fleets and direct guidance on whether to buy new GPUs or free up existing ones, platform teams avoid expensive purchases and long procurement cycles, machine learning teams get capacity faster, and leadership gets better ROI with predictable spend.
- **Accelerate AI delivery:** Stalled workloads are correlated directly to the underlying GPUs, pods and processes running them so that teams can troubleshoot performance bottlenecks in minutes instead of hours, allowing engineers to focus on shipping AI projects.
- **Avoid costly disruptions:** Unhealthy GPUs are proactively identified before failures cascade across a cluster and cause training and inference delays.
- **Maximize ROI on GPU spend:** Teams are empowered and accountable for their GPU utilization and costs, and can easily pinpoint where they are overserving or underutilizing their GPUs. This allows teams to reclaim and reallocate resources in order to reduce wasted spend.

"Datadog GPU Monitoring has made it easy for us to stay on top of our multi-tenant GPU infrastructure. We get per-instance, per-device visibility into core utilization, memory, power and thermals right out of the box with no extra setup. The dashboards are rich out of the gate and simple to customize, and standing up isolated views per customer takes minutes," said Kai Huang, Head of Product at Hyperbolic. "Layering on [LLM Observability](#) ties it all together. We can go from a model latency spike straight to the underlying GPU metrics without switching tools. Full stack AI observability in one platform means both our team and our customers can move faster with confidence."

GPU Monitoring is now generally available. To learn more, please visit: <https://www.datadoghq.com/blog/datadog-gpu-monitoring/>.

About Datadog

Datadog is the leading AI-powered observability and security platform. Our SaaS platform integrates and automates infrastructure monitoring, application performance monitoring, log management, user experience monitoring, cloud security and many other capabilities to provide unified, real-time observability and security for our customers' entire technology stack. Datadog is used by organizations of all sizes and across a wide range of industries to enable digital transformation and cloud migration, drive collaboration among development, operations, security and business teams, accelerate time to market for applications, reduce time to problem resolution, secure applications and infrastructure, understand user behavior and track key business metrics.

Forward-Looking Statements

This press release may include certain “forward-looking statements” within the meaning of Section 27A of the Securities Act of 1933, as amended, or the Securities Act, and Section 21E of the Securities Exchange Act of 1934, as amended including statements on the benefits of new products and features. These forward-looking statements reflect our current views about our plans, intentions, expectations, strategies and prospects, which are based on the information currently available to us and on assumptions we have made. Actual results may differ materially from those described in the forward-looking statements and are subject to a variety of assumptions, uncertainties, risks and factors that are beyond our control, including those risks detailed under the caption “Risk Factors” and elsewhere in our Securities and Exchange Commission filings and reports, including the Quarterly Report on Form 10-Q filed with the Securities and Exchange Commission on February 18, 2026, as well as future filings and reports by us. Except as required by law, we undertake no duty or obligation to update any forward-looking statements contained in this release as a result of new information, future events, changes in expectations or otherwise.

Contact:

press@datadoghq.com



Source: Datadog, Inc.