



# Materials Project of Berkeley Lab Uses Datadog Cloud Monitoring to Simplify Observability on AWS

Learn how Datadog provides Public Sector organizations with a cost-optimized tool to modernize their infrastructures and scale to meet growing demands.

Thank you for downloading this Datadog resource. Carahsoft is the distributor for Datadog cyber solutions available via NASPO, Texas DIR-TSO-4288, MHEC, and other contract vehicles.

To learn how to take the next step toward acquiring Datadog’s solutions, please check out the following resources and information:



For additional resources:  
[carah.io/DatadogResources](https://carah.io/DatadogResources)



For upcoming events:  
[carah.io/DatadogEvents](https://carah.io/DatadogEvents)



For additional Datadog solutions:  
[carah.io/DatadogSolutions](https://carah.io/DatadogSolutions)



For additional cyber solutions:  
[carah.io/cyber](https://carah.io/cyber)



To set up a meeting:  
[Datadog@carahsoft.com](mailto:Datadog@carahsoft.com)  
(703) 921-4160



To purchase, check out the contract vehicles available for procurement:  
[carah.io/DatadogContracts](https://carah.io/DatadogContracts)

Datadog@carahsoft.com | (703) 921-4160

CASE STUDY

# Materials Project of Berkeley Lab Uses Datadog Cloud Monitoring to Simplify Observability on AWS

Learn how Datadog provides Public Sector organizations with a cost-optimized tool to modernize their infrastructures and scale to meet growing demands.

## ABOUT LAWRENCE BERKELEY NATIONAL LABORATORY

Lawrence Berkeley National Laboratory is a US Department of Energy Office of Science lab managed by the University of California. The lab's Materials Project was founded in 2011 to accelerate materials research and time to market for innovation.



Scientific Research



4,200 Employees



Berkeley, California



AWS



“Datadog gives us confidence that there are no blind spots in our cloud architecture. It’s our one-stop shop for observability.”

**Dr. Patrick Huck**  
Senior Computer Systems Engineer  
Lawrence Berkeley National Laboratory


## WHY DATADOG?


- Cost-optimized cloud monitoring at scale
- Enables data visualization
- Provides insight into service usage and faults
- Dashboard capabilities allow for continuous live monitoring screens in Berkeley Lab’s offices
- Enables engineers to conduct continuous active monitoring of systems and alert appropriate teams to potential problems

## CHALLENGE

The Materials Project, a research initiative at Berkeley Lab supported by the US Department of Energy, wanted to make its materials research more accessible to a continually growing number of users by updating its monolithic website.

## USE CASE

 Infrastructure Monitoring

 Container Monitoring

## KEY RESULTS

**5K → 300K users**

Served globally

**2M API requests**


Handled per day

**99.98%**

Uptime achieved

## Executive Summary

The Materials Project at Lawrence Berkeley National Laboratory (Berkeley Lab) migrated its web and API infrastructure to AWS and improved monitoring by using solutions from AWS Public Sector Partner Datadog. The project calculates fundamental material properties to help scientists design and synthesize materials for their applications. It wanted to have more insight into usage while scaling rapidly to meet user needs. The organization used cloud monitoring tools from Datadog to track the usage of its AWS services, including Amazon VPC and AWS Fargate, and fine-tune its offering.



### Gaining Insight into Compute-Intensive Operations at Scale

The [Materials Project](#), a research initiative at [Berkeley Lab](#) supported by the US Department of Energy, wanted to make its materials research more accessible to a continually growing number of users by updating its monolithic website. Founded in 2011, the Materials Project calculates fundamental material properties and develops a host of open-source workflow and analysis software to accelerate materials design. The Materials Project data and its tools are helping researchers in industry, education, and government labs by functioning as a springboard for their research. The project's computations drastically reduce the time for researchers to invent new materials, saving months or even years of painstaking work.

The Materials Project is a compute-intensive effort, and as it scaled to meet US and global demand, its on-premises, monolithic stack strained to power both user and internal needs. The project also lacked insight into service usage and faults. Because the Materials Project is publicly funded, it needed an affordable solution to go along with the modernization of all aspects of its infrastructure stack for a microservice architecture. The organization turned to Amazon Web Services (AWS) Public Sector Partner Datadog to implement a cost-optimized observability solution at scale. The Materials Project can now support 99.98 percent uptime for more than 300,000 users.

### Searching for a Modern Solution

The Materials Project has become a trusted resource for scientists who rely on its data to speed up research and time to market. The project's research has been cited over 19,000 times, and batteries developed using research from the Materials Project are now sold in supermarkets around the world. However, the project's on-premises infrastructure was built for high-performance computing, not for availability. To introduce microservices and enhance the reliability of its website, portals, and APIs, the organization migrated the entire Materials Project website to AWS and redesigned the infrastructure from scratch. As an additional design goal, the project wanted to make open collaboration with other scientists and researchers possible. "We didn't want all the web development to come from our end—we wanted to open it up so that scientists that have Python experience can write app interfaces and improve the Materials Project too," says Dr. Patrick Huck, senior computer systems engineer for the Materials Project at Berkeley Lab. "That required a fundamental change in our approach to designing and delivering website infrastructure."

As the Materials Project migrated to AWS, it needed visibility into its cloud resources. The project chose to use Datadog, a unified observability solution for cloud-scale applications. Datadog offered an evaluation period, and the Materials Project could spend the time developing new architecture to see what exactly it needed. This process culminated in a product demo in February 2022 that featured live data from the Materials Project website. That demonstration convinced the research project that two Datadog tools met its need for cost-optimized cloud monitoring at scale. The Materials Project used [Datadog Infrastructure Monitoring](#) to achieve complete coverage of its estate and [Datadog Container Monitoring](#) for multidimensional visibility into containerized environments. It has also employed Datadog's dashboard capabilities to set up continuous live monitoring screens in its offices.

### Enhancing Access to a Valuable Public Resource

Datadog now provides comprehensive observability across the Materials Project estate.

"Datadog gives us confidence that there are no blind spots in our cloud architecture," says Dr. Huck. "It's our one-stop shop for observability."

The monitoring solution has made it simple for the Materials Project to quickly identify, diagnose, and respond to issues. The team has set up dashboards to keep track of key metrics such as data transfer, uptime, and compute resources at a glance. These dashboards have been useful in identifying bottlenecks and finding ways for the team to optimize its cloud resources.

The Materials Project also chose Datadog because it saw an opportunity to make its cloud expenses resilient against its continued growth using endpoints in Amazon Virtual Private Cloud (Amazon VPC), which gives users full control over their virtual networking environment. Because the Materials Project connects to Datadog in the same AWS Region, it doesn't need to send its AWS reporting data through the public internet, which would not only incur data transfer costs but also significantly increase latency and reduce reliability.

“Datadog offers an Amazon VPC endpoint that handles data transfer between the Datadog instances that perform cloud monitoring and our AWS server instances—without incurring extra traffic costs—while reducing latency and increasing reliability,” says Dr. Huck.”

“It’s given us peace of mind to know that we can control that aspect of our network architecture.”

As part of its redesign, the project uses [AWS Fargate](#), a serverless compute service for containers, in combination with [Amazon Elastic Container Service](#) (Amazon ECS), a fully managed container orchestration service. “AWS Fargate is crucial for the Materials Project,” says Dr. Huck. “Having serverless compute time that is dynamic enough to scale up and scale down as we need makes deployments simple. A lot of the system administration that we used to do is not a factor anymore.” Using AWS and Datadog, the Materials Project can manage its infrastructure with only a handful of scientists-turned-engineers.

The Materials Project had redesigned its architecture, and in September 2021, it prepared to launch the website in new Regions. It worked alongside AWS experts from [AWS Well-Architected](#), a tool and framework to build better in the cloud, to review its offering. “It was a great experience to review and optimize architecture alongside the AWS team,” says Dr. Huck. “It gave us the confidence to launch globally and made it simple for us to focus our resources on what matters.” In 2022, the Materials Project expanded its website to be available globally. The site receives two million API requests and 500 new users daily, growing from 5,000 to over 300,000 users since 2014. “We provide crucial infrastructure for the material sciences,” says Dr. Huck.

### Continuing to Optimize a Global Service

With its new solutions in place, the Materials Project is confident that it can grow to accommodate nearly any user demand, and it will continue optimizing.

“Right now, we’re serving up to 5 TB of data on AWS every month, and our users are satisfied with the reliability that we offer them,” says Dr. Huck.”

[GET STARTED WITH A FREE TRIAL TODAY >](#)

Images are licensed under Lawrence Berkeley National Laboratory.

© The Regents of the University of California, Lawrence Berkeley National Laboratory