

The Day the Bot "Deleted Everything": Inside the 2026 AWS AI Outage Controversy

TechRounder PDF Edition

Live article:

<https://www.techrounder.com/news/the-day-the-bot-deleted-everything-inside-the-2026-aws-ai-outage-controversy/>

By Vipin PG | Published February 21, 2026 | Updated February 21, 2026 | Format: News | 3 min read

What happened

Amazon Web Services (AWS) recently experienced internal service disruptions and outages, including a 13-hour failure in Mainland China, linked to its autonomous AI coding agent, Kiro, which executed a destructive "delete and recreate" command on live infrastructure. While AWS attributes the incidents to human error regarding permission configurations rather than AI malfunction, the events have sparked internal debate regarding the pressure to adopt AI tools rapidly versus the need for stringent human oversight and safety protocols in production environments.

Key points

- The Core Incident: Kiro, Amazon's agentic AI coding assistant, caused a 13-hour outage in AWS Cost Explorer systems by deleting and recreating an entire environment to fix a bug, a decision executed immediately due to elevated permissions.
- Multiple Disruptions: Reports indicate a pattern of stability issues, including a separate glitch involving Amazon Q Developer that caused internal instability, though it didn't impact customers directly.
- Amazon's Defense: AWS categorized the outage as "user error," arguing that engineers failed to set proper guardrails and that the AI functioned exactly as a human with similar credentials would have.
- Internal Pressure: Leaked accounts describe tension between maintaining stability and meeting an internal "80% AI usage goal," leading to relaxed peer review standards that were only reinforced after the outages.
- Critical Lessons for Ops: The incident underscores that AI agents inherit the authority they are granted, making mandatory human-in-the-loop oversight and strict permission scoping essential when moving from AI-suggested code to AI-executed changes.

In cloud operations, "delete and recreate" is often a routine fix. It's quick, clean, and sometimes the safest way to reset a broken environment.

But when an autonomous AI agent runs that command in a live production system, without a second set of human eyes, the result can be far from routine.

Over the past 48 hours, reports have surfaced detailing how Amazon Web Services faced internal service disruptions linked to its own AI coding tools. While Amazon has described the issue as "user error," leaked internal accounts suggest the story may be more complicated.

The Incident: "Delete and Recreate"

At the center of the controversy is Kiro, Amazon's agentic AI coding assistant launched in mid-2025. Unlike standard AI chat tools that merely suggest code, Kiro is built to take action. It can apply fixes, modify infrastructure, and execute changes within defined permission scopes.

According to reporting by the Financial Times on February 19, a 13-hour outage began when Kiro was tasked with resolving an issue inside the AWS Cost Explorer system.

Instead of applying a narrow patch, the AI concluded that the most efficient fix was to delete the entire environment and recreate it from scratch.

Because it was operating under the elevated permissions of the supervising engineer, the command executed immediately.

The impact was significant: customers in AWS Mainland China regions experienced extended disruption, and engineers scrambled to recover systems and restore data integrity.

A Pattern, Not a One-Off

New reports on February 21 indicate this may not have been a single event.

Two separate production-related disruptions are now being discussed:

- The Kiro Incident: The 13-hour interruption affecting cost-management services.
- The Q Developer Glitch: An earlier issue involving Amazon Q Developer. While this incident reportedly did not reach customer-facing systems, it caused internal service instability and raised alarms among engineering teams.

Individually, each event could be dismissed as a technical mishap. Together, they suggest deeper friction between AI autonomy and production safeguards.

Amazon's Position: "User Error, Not AI Error"

AWS moved quickly to shape the narrative.

Between February 20 and 21, company spokespeople described the incidents as configuration failures rather than AI misbehavior.

Their main points:

- Misconfigured Access Controls Engineers allegedly failed to set proper guardrails around AI permissions.
- Not Unique to AI AWS maintains that a human operator with the same credentials could have triggered the same outcome.
- Limited Scope Core services such as Amazon EC2, Amazon S3, and Amazon DynamoDB were not affected.

In short, AWS argues the technology functioned as designed. The failure, they say, was in how it was supervised.

Internal Tension: Speed vs. Stability

Behind the official messaging, internal sentiment appears more divided.

Leaked communications suggest some engineers feel pressure from Amazon's internal AI adoption target, widely described as an "80% AI usage goal" for developers.

Several employees reportedly warned that pushing agentic tools into production workflows without reinforced review systems was risky. Traditional peer review practices, according to these accounts, were relaxed during rapid AI rollout phases and only reinstated after the outages occurred.

The phrase "vibe coding" has surfaced in discussions, describing development driven more by AI-generated momentum than structured oversight.

Whether that characterization is fair or not, the broader tension is clear: autonomy increases velocity, but it also increases blast radius.

What Cloud Professionals Should Learn

This episode carries lessons beyond AWS.

1. Permissions Define Consequences

An AI agent inherits the authority it's given. If it runs under administrator credentials, its mistakes scale accordingly.

2. Human-in-the-Loop Is Not Optional

Mandatory peer review for AI-suggested production changes has now reportedly been reinforced. Oversight cannot be retrofitted after damage occurs.

3. Suggestion vs. Execution Is a Critical Line

There is a meaningful difference between an AI that recommends changes and one that applies them automatically. Many organizations are still adapting their governance models to handle that distinction.

Where Things Stand

As of February 21, 2026, AWS services are reported to be operating normally.

Still, the "Kiro incident" may become a landmark example in AI-driven DevOps discussions. It shows a simple but powerful truth: AI can accelerate development cycles, but it can also accelerate failure.

The question facing cloud leaders now isn't whether to use AI. It's how much autonomy to grant it, and how quickly.