

When Amazon's cloud storage fails, lots of people get wet

March 1 2017, by Mae Anderson



This Sept. 6, 2012, file photo, shows the Amazon logo in Santa Monica, Calif. Amazon's cloud-computing service Amazon Web Services experienced problems in its eastern U.S. region, Tuesday, Feb. 28, 2017, causing widespread problems for thousands of websites and apps. (AP Photo/Reed Saxon, File)

Usually people don't notice the "cloud"—unless, that is, it turns into a massive storm. Which was the case Tuesday when Amazon's huge cloud-computing service suffered a major outage.

Amazon Web Services, by far the world's largest provider of internet-based computing services, suffered an unspecified breakdown in its eastern U.S. region starting about midday Tuesday. The result: unprecedented and widespread performance problems for thousands of websites and apps.

While few services went down completely, thousands, if not tens of thousands, of companies had trouble with features ranging from file sharing to webfeeds to loading any type of data from Amazon's "simple storage service," known as S3. Amazon services began returning around 4 p.m. EST, and an hour later the company noted on its service site that S3 was fully recovered and "operating normally."

THE CONCENTRATED CLOUD

The breakdown shows the risks of depending heavily on a few big companies for cloud computing. Amazon's service is significantly larger by revenue than any of its nearest rivals—Microsoft's Azure, Google's Cloud Platform and IBM, according to Forrester Research.

With so few large providers, any outage can have a disproportionate effect. But some analysts argue that the Amazon outage doesn't prove there's a problem with cloud computing—it just highlights how reliable the cloud normally is.

The outage, said Forrester analyst Dave Bartoletti, shouldn't cause companies to assume "the cloud is dangerous."

Amazon's problems began when one S3 region based in Virginia began to experience what the company called "increased error rates." In a statement, Amazon said as of 4 p.m. EST it was still experiencing errors that were "impacting various AWS services."

"We are working hard at repairing S3, believe we understand root cause, and are working on implementing what we believe will remediate the issue," the company said.

WHY S3 MATTERS

Amazon S3 stores files and data for companies on remote servers. Amazon started offering it in 2006, and it's used for everything from building websites and apps to storing images, customer data and commercial transactions.

"Anything you can think about storing in the most cost-effective way possible," is how Rich Mogull, CEO of data security firm Securosis, puts it.

Since Amazon hasn't said exactly what is happening yet, it's hard to know just how serious the outage is. "We do know it's bad," Mogull said. "We just don't know how bad."

At S3 customers, the problem affected both "front-end" operations—meaning the websites and apps that users see—and back-end data processing that takes place out of sight. Some smaller online services, such as Trello, Scribd and IFTTT, appeared to be down for a while, although all have since recovered.

The corporate message service Slack, by contrast, stayed up, although it reported "degraded service " for some features. Users reported that [file sharing](#) in particular appeared to freeze up.

The Associated Press' own photos, webfeeds and other online services were also affected.

TECHNICAL KNOCKOUTAGE

Major cloud-computing outages don't occur very often—perhaps every year or two—but they do happen. In 2015, Amazon's DynamoDB service, a cloud-based database, had problems that affected companies like Netflix and Medium. But usually providers have workarounds that can get things working again quickly.

"What's really surprising to me is that there's no fallback—usually there is some sort of backup plan to move data over, and it will be made available within a few minutes," said Patrick Moorhead, an analyst at Moor Insights & Strategy.

AFTEREFFECTS

Forrester's Bartoletti said the problems on Tuesday could lead to some Amazon customers storing their data on Amazon's servers in more than one location, or even shifting to other providers.

"A lot more large companies could look at their application architecture and ask 'how could we have insulated ourselves a little bit more,'" he said. But he added, "I don't think it fundamentally changes how incredibly reliable the S3 service has been."

© 2017 The Associated Press. All rights reserved.

Citation: When Amazon's cloud storage fails, lots of people get wet (2017, March 1) retrieved 25 April 2024 from <https://phys.org/news/2017-03-amazon-cloud-storage-lots-people.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.