

Amazon cloud outage: The situation is catastrophic but not serious

By Ahmar Abbas, ITworld | April 22, 2011, 4:37 PM

The recent outage of Amazon Web Services (AWS) east region cloud has taken on many dramatic monikers such as "cloudgate", "cloudburst" and even triggered a creative commiserative competition. Most of us though are not surprised that an outage occurred, but remain a bit puzzled by the length of time it has taken for the engineers to right the situation. We'll look forward to post-mortem reports from AWS that will hopefully guide us to understanding what actually happened. Was there an elusive heisenbug that sprinkled some corrosive pixie dust on the block storage devices? Or was it simply the case of someone making like an air traffic controller and falling asleep at the switch? In any case, full transparency should be the modus operandi here.

Two main themes though quickly emerge out of this episode.

First, is that there are heck of a lot of enterprises out there that are using the public cloud today, and that they have selected the AWS cloud to run their applications. These companies not only are the usual social | local | mobile suspects, but also include companies across media, technology and government sectors. This clear and vigorous adoption of cloud computing now seems to justify the buzz and hype that "cloud" has garnered over the last few years. How else to account for a failure of block storage devices in one of the clouds of one of the cloud providers yielding coverage in CNN, Wall Street Journal and hundreds of other media outlets.

The second theme that sadly emerges is that while a huge number of companies have adopted the public cloud paradigm, their thought process behind design and deployment of their applications on public clouds still seems to follow the traditional datacenter deployment model.

The tremendous ease and benefits of the "programmable cloud infrastructure" that allows a call to an API to set up infrastructure, configure firewalls, provision storage, enable backups and deploy applications in the cloud are not being utilized to automate recovery in case of such catastrophic failures. This becomes all the more painful when you realize that there is minimal incremental cost to have these automations in place. In the public cloud model, companies do not incur reservation costs for their entire recovery infrastructure.

Organizations that leverage native AWS capabilities, such as creating Amazon Machine Images (AMI) for all applications, utilizing snapshots and leveraging one or more of the other 4 geographically isolated AWS regions, can successfully weather these outages. Sure, there will be nuances across the application set and some may not be able to recover gracefully with pure automation and will require manual recovery steps.

Netflix, a large AWS user has institutionalized this in their deployment model. In fact they frequently let loose their Chaos Monkey that constantly forces random failures of even stable AWS instances to ensure recovery. Unlike

Foursquare, Quora and Hootsuite, Netflix did not report any failures during the current AWS east region outage. Recovery.gov a prominent federal government website running on AWS, also recovered quickly and gracefully in another AWS region.

So while the failures have been catastrophic, perhaps embarrassing and will hopefully prompt a review of application deployment and recovery strategies, they are not serious enough to change the dynamics of cloud adoption in short or long term. The benefits of on-demand cloud infrastructure -- such as rapid cycle time, lower capital costs and utility pricing models -- remain strong cloud drivers today, just as they were last week.

Ahmar Abbas is SVP of Cloud Services at San Jose, CA based CSS Corp. CSS Corp is an AWS Solution Provider - so not an entirely disinterested party. Ahmar thanks uber enterprise architect Omar Malick's Facebook status and author Orrin Merton's blog post for inspiring the title of this article.

Source:http://bit.ly/itworld_csscorp