

Information Age

23 May 2011

Black swans in the cloud

An outage at Amazon Web Services and a catastrophic data breach at Sony show that when Internet-scale systems fail, they do so spectacularly

The increasing use of the Internet as a platform and delivery mechanism for computing services appears to be an inexorable march. But two recent incidents gave pause for thought, and revealed that when Internet-scale systems fail it can be catastrophic.

One of the central tenets of cloud computing is that a distributed system is less vulnerable to failure than one that relies on a single piece of hardware. And yet in April 2011 many Amazon Web Services (AWS) customers, including some high-profile websites, lost access to their systems following an outage at one of the company's North American **data** centres.

The issue took three days to resolve, and at the end of it some customers were told that their data was lost forever.

This is precisely the kind of outage that was not supposed to happen in the cloud – Amazon's highly distributed architecture and multiple **data centres** were meant to provide sufficient redundancy as to reduce the risk to zero. It can therefore arguably be described as a 'black swan', an occurrence that is rare, but all the more disruptive when it does happen as a result.

The same might also be said of the multiple data breaches that struck Sony's online gaming services, which saw hackers get their hands on the private data of a staggering 100 million customers.

Cyber attacks are by no means a rarity, but for an electronics company with Sony's reputation to be compromised so effectively was shocking. And the implications are huge, not least for its customers.

Whether or not these events discourage adoption of business cloud services like AWS or consumer cloud services such as Sony's PlayStation Network, they are a reminder that handing data to a third party always involves a certain degree of risk, both for organisations and individuals. And they reveal that a provider's track record is not an infallible measure of what that risk might be.



Alan Calder, CEO of security advisory IT Governance, says the Amazon outage and Sony breach prove that businesses and individuals need to take responsibility for their own data

These two events are proof that there are black swans in technology, just like there are black swans virtually everywhere else.

Both cases prove that whenever you hand your data to someone, you need to ask, "How safe is this?" And the fact that something hasn't gone wrong in the past is no guarantee that it won't happen in the future. Time and time again, people assume that they can take risks without assessing the potential outcome.

The other thing that the two events have in common is the inadequate incident response processes on the part of Sony and Amazon. In both cases, the response was muted and, to one extent or another, fed through lawyers.



Simon Wardley, researcher at CSC's Executive Leadership Forum, says that cloud computing need not be susceptible to black swans

To take full advantage of the cloud, you need to design for failure at every level – not just at the virtual machine level.

The solution to the risk of provider failure is a competitive marketplace of providers offering functionally equivalent services, with easy switching and semantic interoperability between them.

In practice, however, these markets require a common, open source reference model and the first major attempt to achieve this, OpenStack, has only recently begun.

A combination of a marketplace of utility service providers, good enough components and designing systems for failure will create levels of resilience at a given price point that are unobtainable today. This will reduce the likelihood of such black swans far further.