

# Making elastic cloud computing a reality

April 8 2014, by Ry Crozier

---

(Phys.org) —University of New South Wales researchers are using artificial intelligence to create a computer network capable of regulating its own consumption of public cloud services.

The team, led by School of Computer Science and Engineering Lecturer Srikumar Venugopal, hopes the research will make truly elastic [cloud computing](#) a reality.

Elasticity is seen as a key benefit of the cloud, allowing companies to automatically and immediately dial up or down the consumption of cloud computing power based on need.

The idea is to prevent companies over- or under-provisioning resources to run a particular workload.

Presently, companies set rules to manage when to spin-up new virtual servers or shut them down. Systems administrators use their experience and historical data to set the rules.

In future, researchers like Venugopal hope those decisions can be made in an autonomic fashion by the [computer network](#).

His research team has built a software controller that incorporates a simplified version of reinforcement learning, a well-known [artificial intelligence](#) method more commonly associated with robotics than IT.

Under the proposed model, every virtual server instance in the cloud

would have its own controller that monitored the performance of applications hosted on the server.

If an application's performance became critical—for example, too many people visited a website at once—the controller could communicate with others in the network and automatically determine how and where to source extra capacity to cope with the sudden increase in demand.

"These controllers are learning together," Venugopal explains. "They're all doing their own thing but they're talking to each other and then they're figuring out which controller has high load and which one has much less load, and how to balance that out."

The research is ongoing.

"In the latest results, we have incorporated much more resilience—such as workarounds if the virtual machines stop talking to each other," Venugopal says. "We've also improved the intelligence."

The researchers must resolve some outstanding configuration management challenges before the architecture is ready for real-world implementation. These include how to stop one incorrectly configured virtual machine from communicating incorrect information to others, thus causing incorrect decisions to be made.

In parallel, Venugopal's team is examining the potential to elasticise cloud-hosted databases. The Smart Services Cooperative Research Centre funds the research, and the first results were unveiled at a middleware conference held in Beijing in December 2013.

Provided by University of New South Wales

Citation: Making elastic cloud computing a reality (2014, April 8) retrieved 20 March 2024 from <https://phys.org/news/2014-04-elastic-cloud-reality.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.