

# A nimbus rises in the world of cloud computing

May 8 2009

---



Nimbus shows that cloud computing's potential is being realized today. Credit: © 2009 Jupiter Images Corporation

Cloud computing is a hot topic in the technology world these days. Even if you're not a tech-phile, chances are if you've watched a lot of television or skimmed a business magazine, you've heard someone talking about cloud computing as the way of the future. While it's difficult to predict the future, a cloud computing infrastructure project developed at Argonne National Lab, called Nimbus, is demonstrating that cloud computing's potential is being realized now.

So what exactly is [cloud computing](#)? There are varying definitions, but

cloud computing is essentially a form of distributed computing that allows users the ability to tap into a vast network of computing resources through the Internet to complete their work. If, for example, someone wanted to analyze [traffic](#) patterns on the nation's highways, they could upload and store their data into the 'cloud' and have multiple computers crunch the data and then present the results back to them in a unified way as if the work were completed by one giant machine.

Why the word 'cloud'? Some sources believe the term originated in 20th Century telephone systems. Kate Keahey, the lead on the Nimbus project at Argonne, believes the phrase was created when researchers were trying to visualize this type of computing on a whiteboard and made a circular set of squiggles to represent the many components in the internet that would do the computational work. Since these drawings looked liked clouds, Keahey says, researchers soon started saying that data would go 'up to the cloud' for processing.

If all of this sounds familiar, you may have heard this concept before, according to Keahey. Previous decades brought us something called grid computing, which was another type of distributed computing that allowed users to tap into computing resources through a network to get their computational jobs done. But Keahey argues that cloud computing is an evolution of grid computing, with some important differences. With grid computing, you submit what you want computed to a batch scheduler, which puts your job in a queue for a specific set of computing resources, for example a supercomputer, to work on.

"This means you have no control over when your job might execute," Keahey says. You may have to wait as long as a few days before your job is called up, and you're pretty much at the mercy of how that particular grid asset is set up. If its configuration doesn't quite match the complexities of your job, fixing the problem may get very complicated.

Cloud computing, on the other hand, can greatly mitigate this one-size-must-fit-all approach to distributed computing. Many cloud computing platforms allow users to know ahead of time how much computing capacity is available from the cloud, so the work can be done faster. Users can also configure a 'virtual machine' that exists within the cloud to meet the particulars of the jobs they are trying to accomplish. Once a user has configured the type of virtual machine they need for their work, they can go to different cloud computing providers and recreate the system they need to get their jobs done, making computation power a commodity.

Nimbus is an example of such an adaptable system. Keahey and her team developed this open source cloud computing infrastructure to allow scientists working on data-intensive research projects to be able to use such virtual machines with a cloud provider. Nimbus also allows users to create multiple virtual machines to complete specific computational jobs that can be deployed throughout the cloud and still work in tandem with each other. This flexibility allows a user to configure a virtual machine and then connect it to resources on a cloud, regardless of who is providing the cloud.

Having this kind of flexibility and on-demand computing power is vital to projects that are extremely data-intensive, such as research efforts in experimental and theoretical physics. Nimbus has already been deployed successfully to support the STAR nuclear physics experiment at Brookhaven National Laboratory's Relativistic Heavy-Ion Collider. When researchers there needed to turn the massive amounts of data they had generated into viable simulations for an international conference, they used Nimbus to create virtual machines that were run through commercial cloud computing providers.

Creating the virtual machines was relatively easy. "With Nimbus, a virtual cluster can be online in minutes," Keahey says, and the computing

cloud they tapped into provided the computational power they needed to get the simulations done on time. Keahey and her team are now collaborating with CERN in Europe to process the data generated by physics experiments being done there.

Keahey and others in the field believe that this use of custom-crafted virtual machines there are relatively easy to configure on computing clouds will handle more and more of the heavy computational lifting in the future.

Source: National Science Foundation ([news](#) : [web](#))

Citation: A nimbus rises in the world of cloud computing (2009, May 8) retrieved 18 April 2024 from <https://phys.org/news/2009-05-nimbus-world-cloud.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.