

Enabling a new future for cloud computing

August 21 2014



Apt, an NSF-funded precursor to CloudLab, is a testbed instrument that is adaptable to many different research domains through customizable profiles. Shown here is the computer cluster that provides the main hardware resource for Apt. It is comprised of approximately 200 servers that are located in the University of Utah's Downtown Data Center. The Utah portion of CloudLab will go into the same facility - in fact, right next to the equipment in the photographs. Credit: Chris Coleman, School of Computing, University of Utah

The National Science Foundation (NSF) today announced two \$10 million projects to create cloud computing testbeds—to be called "Chameleon" and "CloudLab"—that will enable the academic research community to develop and experiment with novel cloud architectures and pursue new, architecturally-enabled applications of cloud computing.

Cloud computing refers to the practice of using a network of remote servers to store, manage and process data, rather than a local server or a personal computer. In recent years, [cloud computing](#) has become the dominant method of providing computing infrastructure for Internet services.

While most of the original concepts for cloud computing came from the academic research community, as clouds grew in popularity, industry drove much of the design of their architecture. Today's awards complement industry's efforts and will enable academic researchers to experiment and advance cloud computing architectures that can support a new generation of innovative applications, including real-time and safety-critical applications like those used in medical devices, power grids, and transportation systems.

These new projects, part of the NSF CISE Research Infrastructure: Mid-Scale Infrastructure - NSFCloud program, continue the agency's legacy of supporting cutting-edge networking research infrastructure.

"Just as NSFNet laid some of the foundations for the current Internet, we expect that the NSFCloud program will revolutionize the science and engineering for cloud computing," said Suzi Iacono, acting head of NSF's Directorate for Computer and Information Science and Engineering (CISE). "We are proud to announce support for these two new projects, which build upon existing NSF investments in the Global Environment for Network Innovations (GENI) testbed and promise to provide unique and compelling research opportunities that would otherwise not be available to the academic community."

Chameleon

The first of the NSFCloud projects will support the design, deployment and initial operation of "Chameleon," a large-scale, reconfigurable

experimental environment for cloud research, co-located at the University of Chicago and The University of Texas at Austin.



Stampede is one of the most powerful machines in the world for open science research. Funded by the National Science Foundation and built in partnership with Intel, Dell and Mellanox, Stampede went into production on January 7, 2013 at The University of Texas at Austin's Texas Advanced Computing Center (TACC). The Chameleon system will join Stampede at TACC. Credit: Sean Cunningham, TACC

Chameleon will consist of 650 cloud nodes with 5 petabytes of storage. Researchers will be able to configure slices of Chameleon as custom clouds using pre-defined or custom software to test the efficiency and usability of different cloud architectures on a range of problems, from machine learning and adaptive operating systems to climate simulations and flood prediction.

The testbed will allow "bare-metal access"—an alternative to the virtualization technologies currently used to share cloud hardware, allowing for experimentation with new virtualization technologies that

could improve reliability, security and performance.

"Like its namesake, the Chameleon testbed will be able to adapt itself to a wide range of experimental needs, from bare metal reconfiguration to support for ready made clouds," said Kate Keahey, a scientist at the Computation Institute at the University of Chicago and principal investigator for Chameleon. "Furthermore, users will be able to run those experiments on a large scale, critical for big data and big compute research. But we also want to go beyond the facility and create a community where researchers will be able to discuss new ideas, share solutions that others can build on or contribute traces and workloads representative of real life cloud usage."


One aspect that makes Chameleon unique is its support for heterogeneous computer architectures, including low-power processors, general processing units (GPUs) and field-programmable gate arrays (FPGAs), as well as a variety of network interconnects and storage devices. Researchers can mix-and-match hardware, software and networking components and test their performance. This flexibility is expected to benefit many scientific communities, including the growing field of cyber-physical systems, which integrates computation into physical infrastructure. The research team plans to add new capabilities in response to community demand or when innovative new products are released.

Other partners on the Chameleon project (and their primary area of expertise) are: The Ohio State University (high performance interconnects), Northwestern University (networking) and the University of Texas at San Antonio (outreach).

CloudLab

The second NSFCloud project supports the development of "CloudLab,"

a large-scale distributed infrastructure based at the University of Utah, Clemson University and the University of Wisconsin, on top of which researchers will be able to construct many different types of clouds. Each site will have unique hardware, architecture and storage features, and will connect to the others via 100 gigabit-per-second connections on Internet2's advanced platform, supporting OpenFlow (an open standard that enables researchers to run experimental protocols in campus networks) and other software-defined networking technologies.



This is a map of sites that are part of GENI (Global Environment for Network Innovations), a virtual laboratory for networking and distributed systems research and education. Resources available to GENI experimenters include GENI Racks, regional and national backbone networks and WiMAX base stations. GENI and NSFCloud continue NSF's legacy of supporting cutting-edge networking research infrastructure. Credit: GENI Program Office

"Today's clouds are designed with a specific set of technologies 'baked in', meaning some kinds of applications work well in the cloud, and some don't," said Robert Ricci, a research assistant professor of computer science at the University of Utah and principal investigator of

CloudLab. "CloudLab will be a facility where researchers can build their own clouds and experiment with new ideas with complete control, visibility and scientific fidelity. CloudLab will help researchers develop clouds that enable new applications with direct benefit to the public in areas of national priority such as real-time disaster response or the security of private data like medical records."

In total, CloudLab will provide approximately 15,000 processing cores and in excess of 1 petabyte of storage at its three data centers. Each center will comprise different hardware, facilitating additional experimentation. In that capacity, the team is partnering with three vendors: HP, Cisco and Dell to provide diverse, cutting-edge platforms for research. Like Chameleon, CloudLab will feature bare-metal access. Over its lifetime, CloudLab is expected to run dozens of virtual experiments simultaneously and to support thousands of researchers.

Other partners on CloudLab include Raytheon BBN Technologies, the University of Massachusetts Amherst and US Ignite, Inc.

Each team is led by researchers with extensive experience deploying experimental cloud computing systems. Ricci and the CloudLab team have successfully operated Emulab since 2000, providing a network testbed where researchers can develop, debug and evaluate their systems in a wide range of environments. The Chameleon team includes several members of FutureGrid, an NSF-supported testbed that lets researchers experiment in the use and security of grids and clouds.

Ultimately, the goal of the NSFCloud program and the two new projects is to advance the field of cloud computing broadly. The awards announced today are the first step in meeting this goal. They will develop new concepts, methods and technologies to enable infrastructure design and ramp-up activities and will demonstrate the readiness for full-fledged execution. In the second phase of the program, each cloud

resource will become fully staffed and operational, fulfilling the proposed mission of serving as a testbed that is used extensively by the research community.

Provided by National Science Foundation

Citation: Enabling a new future for cloud computing (2014, August 21) retrieved 18 June 2024 from <https://phys.org/news/2014-08-enabling-future-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.