

Envisioning supercomputers of the future

March 18 2016



Researcher Swann Perarnau presents a live demonstration of Argon at the Supercomputing 2015 Conference. Credit: Texas Advanced Computing Center

Last year, President Obama announced the [National Strategic Computing Initiative](#) (NSCI), an executive order to increase research,

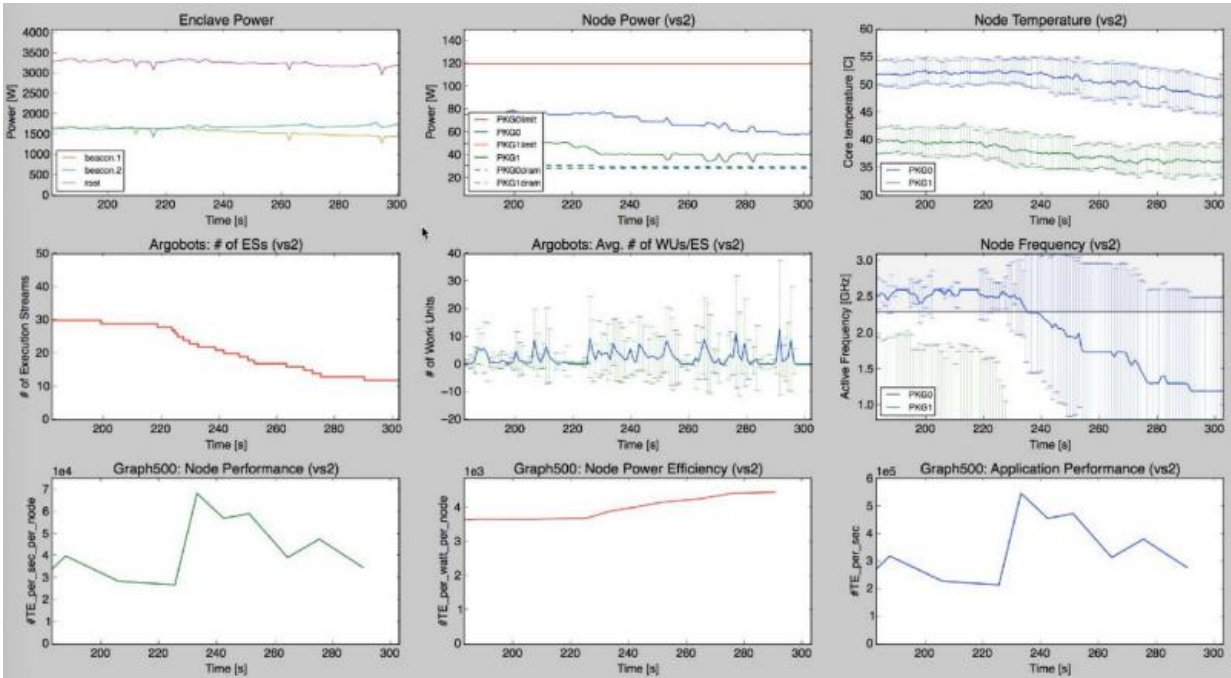
development and deployment of high performance computing (HPC) in the United States, with the the National Science Foundation, the Department of Energy and the Department of Defense as the lead agencies.

One of NSCI's objectives is to accelerate research and development that can lead to future exascale computing systems—computers capable of performing one billion billion calculations per second (also known as an exaflop). Exascale computers will advance research, enhance national security and give the U.S. a competitive economic advantage.

Experts believe simply improving existing technologies and architectures will not get us to exascale levels. Instead, researchers will need to rethink the entire computing paradigm—from power, to memory, to system software—to make exascale systems a reality.

The [Argo Project](#) is a three-year collaborative effort, funded by the Department of Energy, to develop a new approach for extreme-scale system software. The project involves the efforts of 40 researchers from three national laboratories and four universities working to design and prototype an exascale [operating system](#) and the software to make it useful.

To test their new ideas, the research team is using [Chameleon](#), an experimental environment for large-scale [cloud computing](#) research supported by the National Science Foundation and hosted by the University of Chicago and the Texas Advanced Computing Center (TACC).



A live demonstration during SC15 showed Argo's node performance, power efficiency and application performance. Credit: The Argo Project, DOE

Chameleon—funded by a \$10 million award from the [NSFFutureCloud](#) program—is a re-configurable testbed that lets the research community experiment with novel cloud computing architectures and pursue new, architecturally-enabled applications of cloud computing.

"Cloud computing has become a dominant method of providing computing infrastructure for Internet services," said Jack Brassil, a program officer in NSF's division of Computer and Network Systems. "But to design new and innovative compute clouds and the applications they will run, academic researchers need much greater control, diversity and visibility into the hardware and software infrastructure than is available with commercial cloud systems today."

The NSF Future Cloud testbeds provides the types of capabilities Brassil described.

Using Chameleon, the team is testing four key aspects of the future system:

- The Global Operating System, which handles machine configuration, resource allocation and launching applications.
- The Node Operating System, which is based on Linux and provides interfaces for better control of future exascale architectures.
- The concurrency runtime Argobots, a novel infrastructure that efficiently distributes work among computing resources.
- BEACON (the Backplane for Event and Control Notification), a framework that gathers data on system performance and sends it to various controllers to take appropriate action.



Researchers from the Department of Energy-supported Argo Project used Chameleon, a configurable experimental environment for large-scale cloud research, to test a new vision of the operating systems of future exascale computers. Credit: Texas Advanced Computing Center

Chameleon's unique, reconfigurable infrastructure lets researchers bypass some issues that would have come up if the team was running the project on a typical [high-performance computing](#) system.

For instance, developing the Node Operating System requires researchers to change the operating system kernel—the [computer](#) program that controls all the hardware components of a system and allocates them to applications.

"There are not a lot of places where we can do that," said Swann Perarnau, a postdoctoral researcher at Argonne National Laboratory and collaborator on the Argo Project. "HPC machines in production are strictly controlled, and nobody will let us modify such a critical component."

However Chameleon lets scientists modify and control the system from top to bottom, allowing it to support a wide variety of cloud research and methods and architectures not available elsewhere.

"The Argo project didn't have the right hardware nor the manpower to maintain the infrastructure needed for proper integration and testing of the entire software stack," Perarnau added. "While we had full access to a small cluster, I think we saved weeks of additional system setup time, and many hours of maintenance work, switching to Chameleon."

One of the major challenges in reaching exascale is [energy usage and cost](#). During [last year's Supercomputing Conference](#), the researchers demonstrated the ability to dynamically control the power usage of 20 nodes during a live demonstration running on Chameleon.

They released a [paper](#) this week describing their approach to power management for future exascale systems and will present the results at the Twelfth Workshop on High-Performance, Power-Aware Computing (HPPAC'16) in May.

The Argo team is working with industry partners, including Cray, Intel and IBM, to explore which techniques and features would be best suited for the Department of Energy's next supercomputer.

"Argo was founded to design and prototype exascale operating systems and runtime software," Perarnau said. "We believe some of the new techniques and tools we have developed can be tested on petascale

systems and refined for exascale platforms."

More information: Systemwide Power Management with Argo.
[www.mcs.anl.gov/publication/sy ... ower-management-argo](http://www.mcs.anl.gov/publication/systemwide-power-management-argo)

Provided by National Science Foundation

Citation: Envisioning supercomputers of the future (2016, March 18) retrieved 20 April 2024
from <https://phys.org/news/2016-03-envisioning-supercomputers-future.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.