


TECHNOLOGY - Written by [Eric Gedenk](#) on February 10, 2015

New Test Harness Keeps OLCF Resources Running Smoothly

Tags: [HPC](#), [I/O](#), [Lustre](#), [Spider](#), [test harness](#)

SHARE    ...



A test harness is a specific software mix automated to stress the system. OLCF staff created an I/O test harness to catch bugs that happen only at extreme computational scale, such as the conditions surrounding the OLCF's Titan.

Staff collaboration finds that using both synthetic benchmarks and production applications leads to more realistic file system stress testing

High-performance computing ecosystems are complex, including supercomputers, visualization, and analysis clusters to make the scientific workflow appear seamless. High-performance file systems facilitate data sharing among these resources. With each software update or hardware change, though, the system properties change and require testing to ensure continued correctness, functionality, and performance.

Staff at the US Department of Energy's (DOE's) [Oak Ridge National Laboratory](#) (ORNL) saw a need for more efficient systems testing; through a collaboration of groups at the Oak Ridge Leadership Computing Facility (OLCF), a [DOE Office of Science](#) User Facility located at ORNL, they were able to build a state-of-the-art test harness.

A test harness is a specific software mix automated to stress the system. A system's response to this testing can be used to reveal bugs or identify problems that need to be fixed before releasing it to users. In this case, OLCF staff created an I/O test harness to catch bugs that happen only at extreme computational scale, such as the conditions surrounding the OLCF's [Cray XK7 Titan supercomputer](#). The test harness was developed because of Titan's interaction with the Spider II file system, which serves as a data management tool that temporarily stores data as it moves into, out of, and in between the various OLCF computational resources, including Titan.

Sarp Oral, file systems and storage team lead and systems engineer at the OLCF, said that while the OLCF has long used a test harness, earlier approaches were not able to keep up with system complexity. "Until recently, our test harness mostly consisted of synthetic benchmarks," Oral said. "It is resource efficient. You can test a certain aspect of a system using synthetic benchmarks in a much shorter time frame and use fewer resources."

However, according to Oral, the problem with synthetic benchmarks lies in their predefined nature. For example, when cell phone companies perform "drop tests" to prove their products' durability, consumers only see the phone dropped in a very specific way. Generally, though, one does not drop a phone so that it lands

ARCHIVES

[February 2015](#) (4)
[January 2015](#) (3)
[December 2014](#) (5)
[November 2014](#) (11)
[October 2014](#) (5)
[September 2014](#) (3)
[August 2014](#) (5)
[July 2014](#) (5)
[June 2014](#) (8)
[May 2014](#) (3)
[April 2014](#) (4)
[March 2014](#) (5)
[February 2014](#) (3)
[January 2014](#) (3)
[December 2013](#) (5)
[November 2013](#) (13)
[October 2013](#) (2)
[September 2013](#) (5)
[August 2013](#) (9)
[July 2013](#) (4)
[June 2013](#) (9)
[May 2013](#) (2)
[April 2013](#) (2)
[March 2013](#) (3)
[February 2013](#) (3)
[January 2013](#) (3)
[December 2012](#) (5)
[November 2012](#) (5)
[October 2012](#) (6)
[September 2012](#) (5)
[August 2012](#) (4)
[July 2012](#) (8)
[June 2012](#) (4)
[May 2012](#) (5)
[April 2012](#) (2)
[March 2012](#) (4)
[February 2012](#) (9)
[January 2012](#) (1)
[December 2011](#) (8)
[November 2011](#) (7)
[October 2011](#) (4)
[September 2011](#) (4)
[August 2011](#) (8)
[July 2011](#) (6)
[June 2011](#) (8)
[May 2011](#) (3)
[April 2011](#) (4)
[March 2011](#) (2)
[February 2011](#) (7)
[January 2011](#) (1)
[November 2010](#) (2)
[June 2010](#) (1)
[September 2009](#) (1)
[July 2009](#) (1)
[November 2008](#) (1)
[May 2008](#) (1)
[April 2008](#) (1)
[February 2008](#) (1)
[January 2008](#) (1)
[October 2007](#) (2)

CATEGORIES

[Announcements](#) (8)
[Featured](#) (3)
[Headline](#) (1)
[People](#) (95)
[Science](#) (91)
[Support Announcements](#) (2)
[Technology](#) (61)
[Uncategorized](#) (2)
[Video](#) (2)

TAGS

[ADIOS](#) [ALCC](#) [Astrophysics Awards](#) [Big](#)

perfectly on its back as in the cell phone companies' tests. "Things happen," Oral said. "Life is more complex than that, and our systems are complex." Likewise, the increasingly complexity of OLCF systems led staff to look for a more comprehensive test harness.

Despite rigorous testing, when OLCF staff upgraded Spider in early 2014, there were some unanticipated issues. OLCF staff rolled back the Lustre client software on Titan to the original version. Oral said this experience showed the need for a more complete testing procedure and a test harness to avoid these problems in the future. "We brought the High-Performance Computing Operations, Scientific Computing, User Assistance and Outreach, and Technology Integration groups together to discuss how we can better organize ourselves for upcoming tests, how we can put more scientific rigor into our testing, and how we can better test under real-world workloads," Oral said. This integrated team created a comprehensive test harness and improved testing methodologies used at the OLCF, which made it possible for Spider II to be fully updated and implemented in October 2014.

OLCF computational scientist Hai Ah Nam, with the help of the Scientific Computing group, created simplified versions of scientific simulations used on actual production runs for the new test harness. "Our testing methods have improved significantly for two reasons," Nam said. She then explained the success came through the combination of using scientific codes that run regularly on OLCF resources and having an integrated project team to ensure OLCF staff member look for problems from all angles.

Oral reiterated that synthetic benchmarks alone are not capable of adequately stressing the system because of its complexity, size, and unique workload from the users. Using actual code more accurately stress tests OLCF systems, giving OLCF staff better insight into potential problems in the OLCF computing ecosystem.

The next step for OLCF staff is to further refine the test harness by shortening the tests through the use of I/O kernels instead of full applications. By isolating the I/O kernel, OLCF staff can still conduct accurate stress testing, but reduce the testing time, so users see less downtime. "Isolating a kernel, or a shorter snippet of code, is not as easy as it sounds," Oral said. "These codes can be very complicated. You are trying to simplify it to capture the essence of what the code is doing in terms of I/O operations, but you don't want to simplify it so much that you lose the interactions and stress it creates on the system."

Student involvement led to some of the advances, Oral said. "One of our students, Sean McDaniel from the [University of Delaware](#), investigated just this problem last summer, and showed that a publically available I/O kernel did not in fact mimic the I/O from the real application," Oral said. McDaniel's project won the Association for Computing Machinery's Student Research Competition at the SC14 supercomputing conference.

Oak Ridge National Laboratory is supported by the US Department of Energy's Office of Science. The single largest supporter of basic research in the physical sciences in the United States, the Office of Science is working to address some of the most pressing challenges of our time. For more information, please visit science.energy.gov.

[Data](#) [Biology](#) [CASL](#) [Chemistry](#) [Climate](#)
[Combustion](#) [Cray](#) [Data](#) [Director's Discretion](#) [DOE](#) [Early](#)
[Science](#) [Education](#) [Energy](#)
[Engineering Events](#) [File System](#) [Ford](#) [GE](#)
[Global Research](#) [Gordon Bell Prize](#) [GPU](#)
[GPUs](#) [HPC](#) [HPCwire](#) [I/O](#) [INCITE](#)
[Industry](#) [Jaguar](#) [Materials](#)
[Nuclear](#) [Nuclear Physics](#) [NVIDIA](#) [ORNL](#)
[Petascale](#) [Physics](#) [sc13](#) [Spider](#)
[Titan](#) [Training](#) [University of Tennessee](#)
[visualization](#) [Workshops](#)

Home	About OLCF Overview Leadership Team Groups Org Chart Careers Visitor Information & Tours Contact Us	Leadership Science Biological Sciences Chemistry Computer Science Earth Science Engineering Materials Science Physics 2015 INCITE Projects 2014 INCITE Projects 2014 ALCC Projects	Computing Resources Titan Cray XK7 Eos EVEREST Rhea Sith Data Management Visualization and Data Analysis	Center Projects Adios CCI File System Projects IOTA OpenSFS SWTools XGAR	Support Getting Started System User Guides KnowledgeBase Tutorials Training Events My OLCF Software Documents & Forms OLCF Policies	Media Center Center Reports Fact Sheets Media Kit Image Gallery	Summit	SC14
----------------------	------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------	------------------------	----------------------

