Scientific Computing for National Security

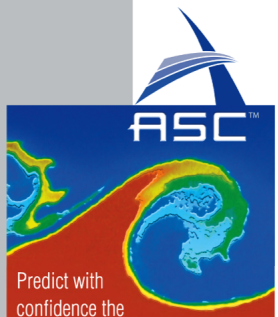

Predict with confidence the behavior of complex systems through comprehensive, science-based simulations

Three labs to leverage common high-performance computing cluster hardware and software

A new Linux cluster at Lawrence Livermore National Laboratory is one prototype of a new tri-lab Linux capacity cluster. Scalable units (SU) will be aggregated into multi-SU clusters of two, four, six, or eight SU, with each cluster available for computing across the three defense laboratories. The multi-SU clusters represent 40, 80, 120, and 160 teraFLOPS resources.

Capacity computing unified across the NNSA complex

For the first time, the Department of Energy/National Nuclear Security Administration (DOE/NNSA) has awarded a single purchase contract that covers all three national defense laboratories—Los Alamos, Sandia, and Livermore. The new computers will provide much needed "capacity" computing, running larger numbers of smaller jobs simultaneously on a single high-performance machine. This allows NNSA's more powerful supercomputers, or "capability" systems, to be dedicated to the larger, more complex calculations critical to the nation's Stockpile Stewardship Program.

In the past, each laboratory conducted its own capacity computing procurements and ran its own operating system and tools. With different hardware and software, these systems provided very different and challenging environments to users, making collaborations difficult and increasing overall costs.

Arctic Ice Concentration

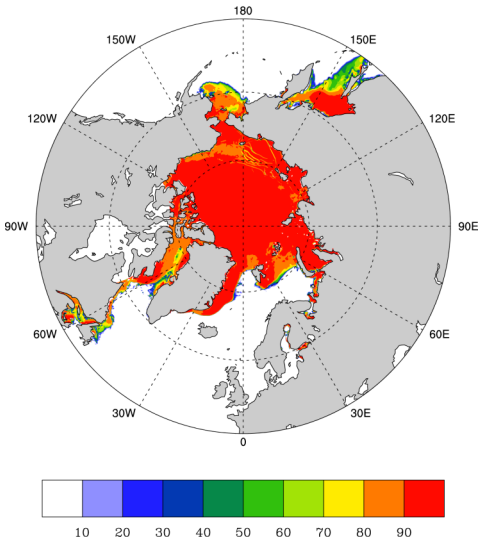

Daily ice concentration (by percent) in the Arctic from a global 0.1 degree coupled ocean and sea-ice model simulation. This simulation was run by Livermore on a prototype Tri-Lab Linux Capacity Cluster (TLCC) machine, using the LANL-developed open-source code POP-CICE.

Advanced Simulation & Computing™

Martin Schoenbauer, NNSA's principal assistant deputy administrator for operations in the Office of Defense Programs, explains the new approach: "Combining the approach for each of NNSA's three laboratories not only saves money but continues to move NNSA towards a smaller, more efficient nuclear weapons complex."

The Tri-Lab Linux Capacity Cluster 2007 (TLCC07) program represents a multimillion dollar subcontract awarded to Appro, a leading developer of high-performance Linux clusters. This subcontract will deliver at least eight scalable Linux capacity clusters to the tri-lab community with an aggregate peak of at least 438 teraFLOPS from at least 21 scalable units (SU). The subcontract also contains options for an additional ten SU for the tri-lab community with an additional aggregate peak of at least 200 teraFLOPS.

Mark Seager, technology lead for Livermore, elaborates, "In the past, high-end Linux clusters were more expensive and difficult to integrate and deploy. By defining this SU 'building block' approach to structure these multiple-cluster deployments using commercially available parts, the labs are able to significantly reduce total costs, simplify the integration and deployment of multiple Linux clusters at all three labs, and put highly needed capacity clusters into production on an accelerated timeline."

Capacity computing

Capacity computing is accomplished through the use of smaller and less expensive high-performance systems to run parallel problems with more modest computational requirements. Capitalizing on the commodity Linux clusters in use throughout the scientific community, ASC has developed a procurement strategy—the Tri-Lab Linux Capacity Cluster (TLCC)—that builds a common hardware environment across the three defense laboratories. The objective is to build, field, and integrate many Linux clusters of various sizes into service through a concept of SU.

With a consistent user environment and common operating system for the capacity computers across the three laboratories, it will be very easy for users to collaborate and run problems at any site. Tripod, a tri-lab committee of computer scientists, is creating robust processes, a roadmap, criteria, and a seamless software environment for the TLCC. The open-source common "software stack" (programs that create a complete computing environment) will be maintained and further developed by researchers at all three laboratories.

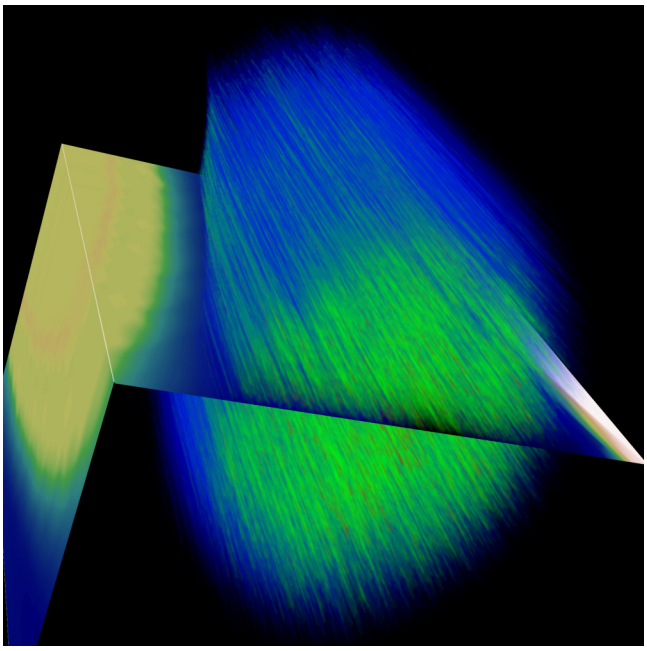

Experiments at the National Ignition Facility (NIF) will use the world's most powerful laser to implode capsules containing deuterium and tritium to the high temperatures and densities at which nuclear fusion occurs. PF3D, run on the prototype Tri-Lab Linux Capacity Cluster (TLCC) configuration, simulates the propagation of a laser beam through ionized gas surrounding the target capsule. Simulations provide crucial input to the design of NIF experiments and greatly reduce the number of laser shots required to achieve ignition.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.
UCRL-BR-235735