

Customer Success Story

Appro Assists LLNL with New Cluster Designed Specifically for Extreme Scale Visualization



"To meet our current and future scientific computing needs requires a visualization cluster with enough memory to generate visualizations of simulation runs on our largest compute platforms as well as sufficient I/O rates for interactive analysis."

Becky Springmeyer

Visualization Group Leader

Advanced Simulation & Computing

Program, LLNL

The Challenge

Lawrence Livermore National Laboratory (LLNL) is home to some of the fastest computers in the world. These machines are solving extreme scale problems using many tens of thousands of processing cores. For example, LLNL recently installed the Dawn cluster which can deliver 500 TFLOP/s of performance. In 2012, LLNL expects to have the Sequoia supercomputing cluster operational with a projected performance of over 20 PFLOP/s.

These systems will be focused on strengthening the foundations of predictive simulation through running very large suites of complex simulations and then comparing model predictions with experimental data. In addition, the machines will be used for weapons science calculations necessary to build more accurate physical models. This work is a cornerstone of NNSA's Stockpile Stewardship program to ensure the safety, security, and reliability of the U.S. nuclear weapons stockpile today and into the future without underground testing.

Such large clusters are capable of crunching through numbers at an alarming rate. While attention is often placed on the FLOP/s rate, the amount of data generated by these systems is equally amazing. It is not uncommon to produce multi-terabyte data sets with tens of billions of zones, thousands of files per time step, and hundreds of time steps. The need to visualize this data is critically important to the program's success. One of the challenges facing the LLNL visualization team was the design and procurement of a cluster with enough horse power to manage the large data sets generated on clusters like Dawn. As stated by Becky Springmeyer, Visualization Group Leader in LLNL's Advanced Simulation and Computing Program;

"To meet our current and future scientific computing needs requires a visualization cluster with enough memory to generate visualizations of simulation runs on our largest compute platforms as well as sufficient I/O rates for interactive analysis,"

An often forgotten aspect of HPC, these visualization resources provide an essential and critical part of the high performance computing environment by providing computational scientists with the ability to interactively explore very large data sets as well as enabling large-scale post-processing and data reduction operations.



Figure 1: LLNL Clusters

The Solution

In order to visualize the large amounts of data, LLNL requested an Appro Hyper Supercomputing Cluster specifically designed to support interactive data analysis on the extreme scale computing systems. The successful Scalable Unit (SU) concept, pioneered by LLNL, LANL, Sandia National Lab (SNL), and Appro, the experienced leader in SU cluster design architecture was chosen as the solution provider. Aptly named "Graph", the Appro Hyper Cluster was deployed at LLNL using four Scalable Units and has a measured performance of 110 TFLOP/s. Similar to previous procurements, the Appro Hyper Cluster was based on the successful combination of Appro AMD Opteron based servers, InfiniBand interconnect, and a shared software stack developed by the Tri-Labs. The ease of procurement and installation leveraged the common hardware and software platform that allowed easy growth of clusters in incremental units while reducing costs.

The compute nodes are based on six-core AMD Opteron™ processors. The cluster includes 564 compute nodes for a total of 2,256 processors (4 per node). Using six-core AMD Opteron™ processors (2.0GHz) Appro was able to deliver a total of 13,536 processing cores. In order to achieve effective visualization performance, each node was given a full complement of 128 GB of RAM for a total of 73 TBytes of memory. InfiniBand was used as the interconnect between nodes. Forty eight Flextronics InfiniBand (20 Gb/s DDR) edge switches, and 2 Voltaire (20 Gb/s DDR) Spine switches were used for the switch fabric.



Figure 2: Graph Cluster on datacenter floor at LLNL

The largest machine at LLNL to be supported by the Appro Hyper Cluster is Dawn. Delivered to the Lawrence Livermore National Laboratory in January and February, Dawn (an IBM Blue Gene/P system) will lay the applications foundation for multi-petaFLOPS (floating point operations per second) computing on Sequoia, a 20-petaFLOPS IBM system to be delivered in late 2011. These supercomputers are capable of running very large suites of complex simulations and the Appro Hyper Cluster is capable of supporting complex visualization and analysis tasks on data sets generated on these much larger machines. These interactive visualization and analysis applications, often called post-processing, are different from those run on the typical computational (number crunching) batch cluster. Interactive exploration of large data sets requires a memory-rich environment with I/O optimized through dedicated file system connections. Post-processing tasks are heavily I/O bound, so specialized visualization servers that optimize I/O rather than CPU speed are better suited for this work. The Appro Hyper Cluster has high-speed connections to the Lustre file system in order to provide users with zero-copy access to their files. Both memory and I/O rates are key enablers of the Appro Hyper Cluster as there needs to be enough memory to generate visualizations of the largest data sets and sufficient I/O for interactive analysis.

The Summary

LLNL has procured a new visualization cluster geared specifically to support interactive data analysis and visualization on the extreme scale. Visualization resources provide an essential and critical part of the high performance computing environment providing computational scientists with the ability to interactively explore very large data sets as well as enabling large-scale post-processing and data reduction operations.

LLNL needed a cluster with much more memory and high-speed connections to Lustre than existing LLNL

visualization clusters. By using Appro's experience with SU cluster design, the Appro Hyper Cluster was quickly procured and brought into service. Users can now experience a reduction in the time required for multi-day visualization problems as well as maintaining a high level of interactivity in analyzing the data sets generated on the largest LLNL computers.

About Lawrence Livermore National Laboratory

Established in 1952, Lawrence Livermore National Laboratory (LLNL) is a premier applied science laboratory, part of the National Nuclear Security Administration within the Department of Energy. LLNL is responsible for ensuring the nation's nuclear weapons remain safe, secure, and reliable through application of advances in science and engineering. With its unique capabilities, the Laboratory meets other pressing national security needs, including countering the proliferation of weapons of mass destruction and strengthening homeland security against the terrorist use of such weapons. The Laboratory is an international leader in many areas of science and technology and undertakes significant research programs in energy, environment, bioscience, biotechnology, and basic science and advanced technology. Since so many of its projects require huge volumes of data analysis and hosting of highly complex proprietary applications the need for more powerful simulation compute resources has been an integral requirement since the 1980's.

About Appro International, Inc.

Appro is a leading developer of supercomputing solutions. Appro is uniquely positioned to support High Performance Computing markets focusing on small to large-scale deployments where lowest total cost of ownership is a primary consideration. Appro accelerates technical applications and business results unlocking the value of IT through outstanding price/performance, balanced architecture, open standards and engineering expertise. Appro headquarters is in Milpitas, CA with offices in Korea and Houston, TX. To learn more go to <http://www.appro.com>

Learn More

To learn more about Appro Supercomputer Solutions please contact your Sales Representative, visit www.appro.com or email us at clustersolutions@appro.com

For more information about Lawrence Livermore National Laboratory visit www.llnl.gov. Information on LLNL Clusters can be found at <https://computing.llnl.gov/linux/clusters.html>.

To learn more about AMD's latest ground-breaking technologies, visit www.amd.com.



Appro International, Inc. | 446 South Abbott Avenue, Milpitas, CA 95035 | Toll Free 800.927.5464 | www.appro.com

Appro International, Inc. All rights reserved. Appro and the Appro logos are registered trademarks of Appro International, Inc. Other company products or service names are the property of their respective owners.